TMOD Technology and Science Program News
June 2000
Issue 12

In This Issue

Charles Stelzried and Michael Klein

In this issue's first article, David Rochblatt and Vic Vilnrotter demonstrate the advantage of combining technologies for optimizing the G/T performance of the 70-m antenna at Ka-Band. A Deformable Flat Plate corrects for systematic deformation, whereas a seven-element Array Feed Compensation System provides real-time compensation for time-varying losses. Next, Yoaz Bar-Sever describes a breakthrough in Global Differential GPS Technology performance with a new software set. Real-Time Net Transfer is expected to enable 10-cm real-time positioning accuracy for airborne remote-sensing platforms, as well as serve as a prototype for the Mars navigation and communications infrastructure network in support of future Mars explorations.

Design compromises determine the optimum antenna configuration for meeting mission requirements. Suzanne Spitz, et al., present examples of engineering variations reflecting these compromises. Inflatable reflectarrays and layered lenses seem optimum for large apertures and a displaced-axis dual-reflector antenna is expected to provide minimum storage volume with high efficiency and relatively small aperture size. Maximizing science data return with data compression and intelligent onboard science processing promises to provide major improvements in optimizing future deep space communications. Aaron Kiely, et al., discuss these concepts and how they differ from commercial off-the-shelf techniques.

The Multi-mission Encrypted Communication System and the Web Interface for Telescience enable Internet-based Mars lander and rover mission operations. Paul Backes and Jeffrey Norris discuss how an increased number of scientists will participate in planetary rover missions at reduced costs using this approach. Scientists from other institutions can participate in operational testing without having to travel to JPL. Mark James, et al., describe an autonomous fault monitoring system using the JPL-developed tools Beacon-based Exception Analysis for Multi-missions (BEAM) and Spacecraft Health Inference Engine (SHINE). BEAM is used as an advanced prognostic state estimator and SHINE is used for real-time diagnostics and interpretation of the output from BEAM.

Sami Asmar and Trish Priest describe the productivity of a newly implemented remote operations capability developed by the Radio Science Systems Group. Initial results from 5,000 occultations of the Mars Global Surveyor spacecraft have yielded 88 vertical profiles of the Martian neutral atmosphere. An exploratory JPL–UCLA experiment using radio astronomy observations to detect ultrahigh energy neutrinos is described by Peter Gorham et al. Theoretical predictions suggest that UHE neutrinos, massless particles with energies a billion-billion times the energy of visible light, can pass through the moon and produce nanosecond pulses of microwave radio emission.

The Goldstone-Apple Valley Radio Telescope science education project is partnering with the Casssini mission to enhance the science return from the spacecraft. Data from student observations will be used to accomplish an in-flight calibration of a microwave radiometer, which will then perform previously unplanned microwave observations.

Demonstration of a Ka-Band Array Feed–Deformable Flat Plate Compensation System

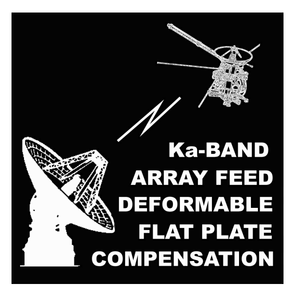

David Rochblatt and Victor Vilnrotter

Introduction

One of NASA's current technology initiatives is to increase its deep space communications capacity by the implementation of Ka-Band (32 GHz) reception on all Deep Space Network (DSN) antennas. By increasing the carrier frequency from 8.4 GHz to 32 GHz (corresponding to X-Band and Ka-Band, respectively), an overall gain of more than 6 dB in communication performance is expected, including atmospheric effects, without increasing spacecraft transmitted power or antenna aperture.

Continued on next page

CONTINUED FROM PAGE 1

However, achieving adequate antenna aperture efficiency at Ka-Band with the large DSN antennas represents a significant challenge. At DSS-14, for example, gravity-induced deformations of the primary antenna surface lead to losses in aperture efficiency at Ka-Band of 3.6 dB at low elevation angles and 6.5 dB at high elevation angles. Systematic deformations, which can be accurately modeled, are best compensated by an open-loop technology, while unmodeled losses and losses due to time-varying distortions and pointing errors induced by wind, thermal gradients, and subreflector misalignments require real-time compensation. Two distinct regions in the antenna RF beam path where these technologies could be implemented are the antenna focal plane and the area where the primary reflector surface is imaged. A seven-element Array Feed Compensation System (AFCS) located in the antenna focal plane would provide real-time compensation for time-varying losses and pointing errors. The Deformable Flat Plate (DFP) compensation system, an open-loop technology that can be used to correct for systematic deformation, could be located in the latter region.

The DFP and AFCS were designed and constructed at the Jet Propulsion Laboratory and have been evaluated at GDSCC for several years [Vilnrotter et al. 1992, Vilnrotter et al. 1995, Bruno et al. 1996]. Demonstrations on the 70-meter antenna at DSS-14, which culminated in the AFCS–DFP experiments, are described in [Rochblatt 1999, Vilnrotter and Fort 1999] and summarized in this article. In this combined configuration, the DFP functioned as a controllable RF mirror placed in front of the AFCS, refocusing divergent RF fields into the array, thereby increasing the real-time recovery of signal energy by the AFCS. Analysis of the data collected from these experiments demonstrates that a joint AFCS–DFP compensation system, working in real-time, could recover most of the signal energy lost to gravitational deformations.

Array Feed Compensation System

The AFCS is comprised of a feed array and cryogenic dewar, Downconverter Assembly, and Digital Signal Processing Assembly (DSPA). The feed array consists of seven identical smooth-walled Potter horns, each 4.4 cm in diameter, providing 22-dBi gain at 32 GHz. The wall of each horn is tapered near the top to facilitate close packing, minimizing the loss of signal energy between horns. The cryogenic front end contains three-stage High Electron Mobility Transistor (HEMT) low-noise amplifiers that provide 25 dBi of total gain, thus establishing high SNR without adding significant receiver noise to the signal. Each HEMT low-noise amplifier outputs an amplified RF signal at 32-GHz center frequency, which is then translated to 300-MHz intermediate frequency (IF) by the seven-channel Downconverter Assembly and transmitted to the DSPA. The DSPA estimates the combining weights in real time, and combines the weighted digital samples to produce a “combined output” with improved SNR.

The DSPA has been designed to process either broadband sources (data-modulated spacecraft carriers and natural radio sources) or narrowband sources (the residual carrier of a Ka-Band spacecraft). Either approach can be selected by the user: In each case the final result is a “combined channel” output with maximized SNR, although different techniques must be employed to estimate the optimum combining weights for the two cases. In addition, the DSPA also extracts pointing information from the complex baseband samples, thus providing the capability for closed-loop antenna tracking based on real-time array feed measurements. During the holography cone experiments, the user could choose either broadband or narrowband versions of an accurate and well-tested AFCS “boresighting” algorithm, or an experimental “least-squares” tracking algorithm that was installed and briefly demonstrated during these experiments.

AFCS Antenna Deformation Compensation

The New Millennium spacecraft, DS1, provided a stable Ka-Band signal throughout the holography cone experiments. During a typical DS1 track, it was necessary to update antenna pointing in real time using the coherent version of the array feed boresighting algorithm [Vilnrotter and Fort 1999]. In addition to providing pointing information, this algorithm records the complex signal voltages at the end of each integration time in all seven channels, as well as in the combined channel, measures rms noise voltages, and updates optimum combining weights. This data can be used to compute the combining gain, G_c, which is defined as the ratio of the SNR of the combined channel to the SNR of the central channel, and is an important quantity for characterizing AFCS performance.

CONTINUED ON PAGE 16

An Internet-Based Global Differential GPS Technology

GDGPS TRACKING

Yoaz Bar-Sever

Overview

Precise, real-time, onboard knowledge of a spacecraft's state (e.g., its position and velocity) is a critical component of intelligent platform control. This article describes a recent technology development that provides a breakthrough in this capability for near-Earth platforms carrying a Global Positioning System (GPS) receiver—seamless global coverage and roughly an order of magnitude improvement in accuracy compared to state-of-the-art.

In a paper presented at the Institute of Navigation National Technical Meeting last January in Anaheim, California, entitled "An Internet-Based Global Differential GPS System, Initial Results," Muellerschoen et al. [2000] (Section 335) described the world's first real-time, ultraprecise, global differential GPS (GDGPS) system. Using novel Internet-based technology for editing and real-time streaming of GPS data from NASA's Global GPS Network (GGN), combined with the real-time orbit-determination flight software, Real-Time GIPSY (RTG), Muellerschoen et al. [2000] reported achieving ~20-cm real-time positioning accuracy for a ground-based GPS receiver. This performance is roughly an order of magnitude better than any differential service currently available. Although a number of private and government organizations provide localized real-time positioning services to users on or near the ground, a system capable of supporting global space users has never been achieved—or even attempted—due to the perceived technical and cost challenges.

This technology will enable us to provide cm-level, real-time, onboard orbit determination for Earth-orbiting satellites, and 10-cm-level, real-time positioning accuracy for airborne remote sensing platforms such as AirSAR. GDGPS will revolutionize our ability to sense and respond in a timely manner to natural hazards such as earthquakes and volcanic eruptions. With ultrahigh-accuracy orbit determination available in real-time, NASA (and other satellite operators) would be able to bypass an entire operational element which is currently required on the ground and consumes significant resources, even for a single orbiter such as TOPEX/Poseidon. The savings in this area alone could easily amount to between hundreds of thousands to millions of dollars per year, per mission. Furthermore, these savings would be multiplied by a large multi-mission factor, since the GDGPS system could support an unlimited number of users simultaneously—the ultimate in multi-mission service. We envision this technology also serving as a prototype for a similar system that would be part of the Mars navigation and communications infrastructure network in support of future Mars explorations.

The development of GDGPS was supported over the last two years by a small TMOD Technology Task, with co-funding from JPL's Technology Applications Program. This technology development effort leverages the significant investment NASA has made in the GGN, as well as the U.S. government's investment in the Wide Area Augmentation System (WAAS), which was developed at JPL.

Dissemination of differential corrections to authorized users is currently enabled by specially developed Internet tools. We are pursuing commercial and institutional support for the development of an effective global communications system that will relay differential corrections to users anywhere, on the ground and in space. We are also planning to modify existing GPS receivers to enable the reception and processing of broadcast corrections. Currently, the correction signal can be combined with GPS data in the flight computer. Suitably modified GPS receivers will be able to receive the correction signal and internally combine it with the GPS data. With the new generation of GPS space receivers in the

CONTINUED ON NEXT PAGE

CONTINUED FROM PAGE 3

GPS-On-A-Chip family (such as JPL's BlackJack receiver), the embedded PowerPC processor has the power to perform significant calculation onboard in real time, as well as handle the tasks which would normally be associated with real-time navigation and control decisions.

System Description

The JPL architecture for a real-time, global, differential GPS system was first put forward by Yunck et al. [1995, 1996]. A commercial North American Wide Area Differential GPS (WADGPS) system based on the JPL architecture and software was implemented in 1997 by SATLOC, Inc. (a commercial company), primarily intended for the agricultural market [Bertiger et al. 1998]. In 1996 the Federal Aviation Administration (FAA) selected JPL's architecture and software for their prototype WAAS, which has been installed in the FAA's National Satellite Testbed (NSTB) and is scheduled to become operational in 2001. WAAS provides GPS differential corrections for users in the continental United States (CONUS). The WAAS and SATLOC systems, like every other existing differential system, are optimized for users carrying a single-frequency GPS receiver. These receivers are susceptible to large, unmodeled ionospheric delays. In order to compensate for this error source, differential systems employ a dense network of reference stations distributed over their service area (for example, WAAS uses 24 reference stations in the continental U.S.), and produce maps of the total electron content (TEC) of the ionosphere, which are transmitted to the users. Residual ionospheric errors and a variety of other error sources yield an overall user positioning accuracy not better than 1-m RMS, and sometimes much larger than that. No differential services cover the entire globe, and none are capable of supporting spaceborne users.

Our new system is geared toward users carrying dual-frequency receivers, such as are flown on a wide variety of remote-sensing, low-Earth-orbiter missions, including: TOPEX/Poseidon, Jason-1, GPS/MET, SAC-C, Champ, IceSat, VCL, SRTM, and others. These high-end users typically require high-accuracy positioning. Having eliminated the ionosphere as an error source using dual-band receivers, these users are still susceptible to errors related to GPS ephemerides and clocks. Ground-based users and aircraft must also deal with errors due to the effects of the troposphere. Accurate correction for GPS ephemeris and clock errors requires a network of GPS reference sites. Zumberge et al. [1997] demonstrated that a well-distributed global network of about a dozen sites is sufficient for continuously providing GPS ephemeris and clock corrections for GPS satellites. We have taken advantage of the NASA GGN, which is operated and maintained by JPL and consists of approximately 50 sites (traditionally operated in batch mode [**http://igscb.jpl.nasa.gov/**]).

The breakthrough in our capability came in the form of a new software set, Real-Time Net Transfer (RTNT), which was designed to return GPS data in real time from remote receivers. RTNT collects, edits, and compresses the raw GPS observables at the remote site. It then transmits the packetized data over the open Internet to the processing center. At the processing center, the global data is analyzed by the real-time orbit-determination software, RTG, to produce precise GPS orbits and clocks. These are formatted as corrections to the GPS broadcast ephemerides, encoded, and provided to authorized users over the Internet. The combined software set was named IGDG (Internet-based Global Differential GPS). The two major software modules are described below.

Real-Time Net Transfer

The remote sites minimally have a dual-frequency GPS receiver, a PC running linux operating system, and connectivity to the open Internet. A receiver-specific data daemon running on a PC at the remote site establishes communications with the receiver through its serial port, and places the raw GPS data (phase and range) in a revolving buffer of shared memory. A second process (that is independent of receiver type) reads this shared memory and opens a socket connection to the central data daemon. The data is checked and edited, and then sent out through the socket. This step is critical, because a certain percentage of the data is lost in the transmission over the Internet (see below), making data editing at the processing center impossible. A remote site tracking 10 GPS satellites transmits data over the open Internet at a rate of 227 bytes/sec to a central data daemon. At the central processing site, the transmitted data are collected by another RTNT process that monitors the state of the whole system. This central data daemon sorts the data according to timetag, rejects duplicate transmissions, and at a specified drop-dead time, outputs all the data at a common epoch into a circular buffer of shared memory.

Reliability is a key consideration of an operational differential service. For improved reliability, the central data daemon keeps track of the sequence number of packets arriving from each remote site, and may request up to three retransmissions of missed data epochs. A unique architecture of fully redundant processing centers guarantees continuous service even if one processing center unexpectedly goes down. The central data daemon has a twin data daemon running on another computer; all of its incoming GPS data is relayed to its twin via socket communications. Should the twin no longer "see" any data flow, it will send out a request to the entire global network for re-routing the real-time data to itself. It would then serve as the central data daemon until the primary daemon is brought back on-line. Another possible approach is to chain these data daemons in order to export the real-time GPS data to any other computer on the open Internet, and even merge streams from various data daemons or additional receivers.

The global differential corrections produced by RTG are packaged into a 560 bit/sec message, and can be made available on the open Internet via a TCP server running at JPL.

Real-Time GIPSY

RTG provides one-second GPS clocks and real-time estimates of the dynamic GPS orbits [Bertiger et al. 1998]. RTG contains many of the precise models of the GIPSY OASIS II (GOA II) software. GOA II has a long history of precise orbit determination of GPS and other spacecraft, and precise GPS geodetic applications. Post-processing of global GPS data with GOA II routinely yields better than 10-cm GPS orbits.

Orbit estimates are needed less frequently than the clocks due to their slower varying physical behavior. RTG reads the shared memory output of the central data daemon process. Orbit and troposphere estimates at the reference stations are computed once per minute by RTG. These corrections are then placed into another revolving buffer of shared memory so that they may be read by the clock correction process, which produces corrections at 1 Hz.

RTG is also used for onboard, autonomous user positioning. It has been embedded in real-time user equipment for flight on the X33 sub-orbital vehicle, and has flown on the NASA DC-8 SAR flights [Muellerschoen and Bertiger 1999]. In this mode, RTG ingests the correction message as well as the raw GPS data from a given receiver and provides precise estimates of the user position. For users with known dynamics (such as spacecraft), it performs orbit determination. For users with unknown dynamics (vehicles, airplanes), it provides kinematic positioning.

Results

Currently, real-time data is returned to JPL from 15 out of the 50 GGN sites (Fig. 1). This network does not provide 100 percent GPS

CONTINUED ON NEXT PAGE

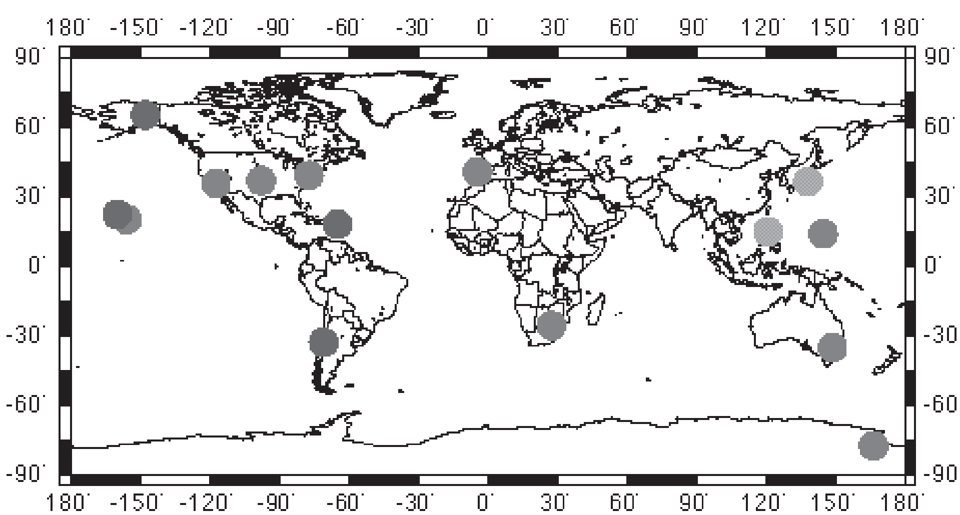

Figure 1. Network of dual-frequency GPS receivers returning data to JPL in real time

CONTINUED FROM PAGE 5

coverage and, as a result, GPS orbit determination accuracy is somewhat degraded. GPS data at 1 Hz are returned with a latency of <1.5 seconds using the open Internet. (Data latency is the difference between the GPS timetag of the data and the central computer's system time.) Better than 95 percent of the data are returned in less than 3 seconds (Fig. 2). GPS orbits are determined to an accuracy of about 50-cm RMS, relative to the precise orbits produced by GOA II in post-processing (Fig. 3). Orbit and clock estimates are then differenced with the broadcast ephemeris and clock. The corrections, which have been produced at JPL in test mode since November 1999, can be automatically distributed via the Internet to authorized users.

To test the accuracy of these global differential GPS correctors, position-stationary GPS receivers are pointed at known locations. Kinematic positioning is insensitive to the dynamics of the user. Here RTG is used with global differential GPS correctors obtained via a socket connection from the TCP server. Comparing the kinematic positioning of a static user with its post-processed location shows a 19-cm RMS vertical error (Fig. 4). Most of the positioning error is believed to be due to GPS orbital error caused by incomplete net-work coverage. GPS orbit errors are expected to improve as more GGN sites are upgraded to provide data in real time. Indeed, this performance constitutes a significant improvement over the 40-cm RMS error in AirSAR positioning

Figure 2. Typical real-time GPS data latencies

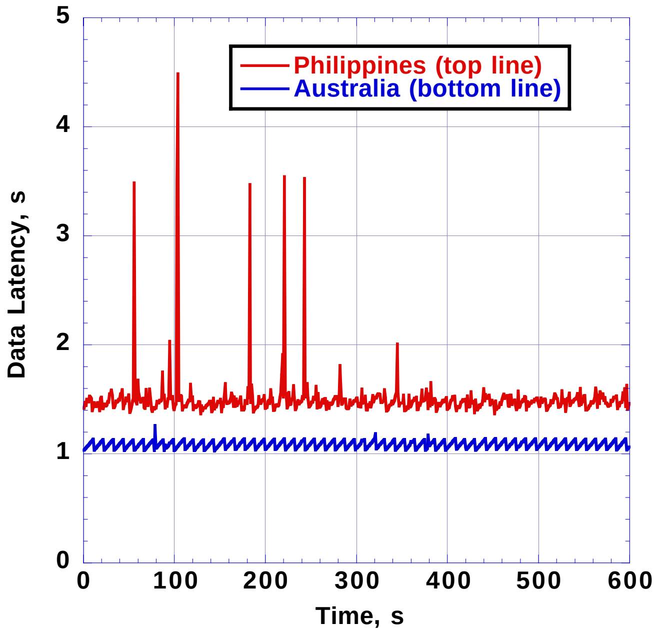

Figure 3. RSS 3D orbit errors for GPS22 for a 24-hr period with RTG and a 15-station global network

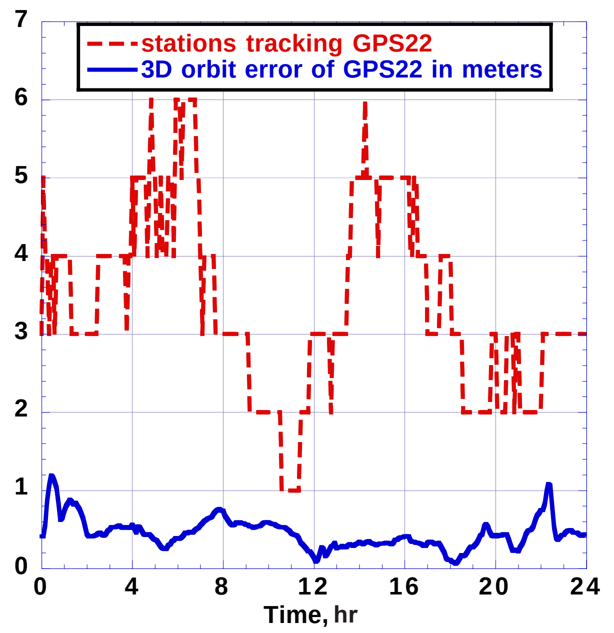

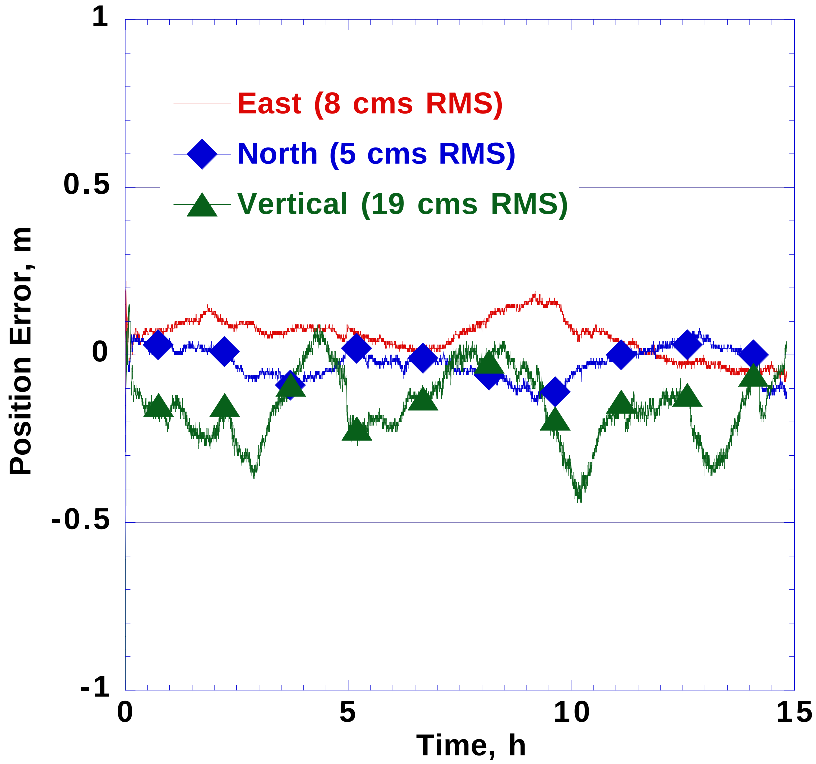

Figure 4. Results of 15 hours of 1-Hz positioning of a static receiver at Fort Pierce, Florida, with global diffferential corrections

using differential corrections based on SATLOC's CONUS-only network [Muellerschoen and Bertiger 1999]. Consequently, we anticipate 10-cm positioning accuracy with a network of 20 sites.

Future Plans

The accuracy and coverage of the GDGPS system is currently without parallel. However, many improvements remain to be realized. Most important is the densification of the real-time network to provide the redundancy that is required by an Internet-based architecture. Various components of the system require tuning for optimal performance in specific configurations. We are working to optimize the real-time GPS orbit-determination strategy as the size of the real-time network grows, and to optimize the user positioning strategy for specific applications, such as orbit determination, airborne, or terrestrial.

While the technology to generate global differential corrections has been demonstrated, it falls short of a differential service because it lacks two critical components, namely, a global communication relay system and user equipment. Our current differential communications technology is completely Internet-based. We are now developing the missing components in order to provide NASA with a revolutionary and fully operational new capability.

References

Bertiger, W. I., Y. E. Bar-Sever, B. J. Haines, B. A. Iijima, S. M. Lichten, U. J. Lindqwister, A. J. Mannucci, R. J. Muellerschoen, T. N. Munson, A. W. Moore, L. J. Romans, B. D. Wilson, S. C. Wu, T. P. Yunck, G. Piesinger, and M. Whitehead, "A Real-Time Wide Area Differential GPS System," *Navigation: Journal of the Institute of Navigation*, Vol. 44, No. 4, 433–447, 1998.

Muellerschoen, R. J., W. I. Bertiger, "Flight Tests Demonstrate Sub 50 cms RMS. Vertical WADGPS Positioning," Proceedings of ION GPS-99, Nashville, Tennessee, 1999.

Muellerschoen, R. J., W. I. Bertiger, M. F. Lough, D. Stowers, and D. Dong, "An internet-based global differential GPS system—initial results," Proceedings of the ION National Technical Meeting, Anaheim, California, 2000.

Yunck, T. P., W. I. Bertiger, S. M. Lichten, A. J. Mannucci, R. J. Muellerschoen, and S. C. Wu, "A Robust and Efficient New Approach to Real Time Wide Area Differential GPS Navigation for Civil Aviation," JPL Report D-12584, 1995.

Yunck, T. P., Y. E. Bar-Sever, W. I. Bertiger, B. A. Iijima, S. M. Lichten, U. J. Lindqwister, A. J. Mannucci, R. J. Muellerschoen, T. N. Munson, L. Romans, and S. C. Wu, "A Prototype WADGPS System for Real Time Sub-Meter Positioning Worldwide," Proceedings of ION GPS-96, Kansas City, Kansas, 1996.

Zumberge, J. F., M. B. Heflin, D. C. Jefferson, M. M. Watkins, and F. H. Webb, "Precise Point Positioning for the Efficient and Robust Analysis of GPS Data From Large Networks," *J. Geophys. Res.*, 102, 5005–5017, 1997.

Technology Roadmap for High-Gain Spacecraft Antennas

Suzanne Spitz, Aluizio Prata, John Huang, and Luis Amaro

Introduction

It is a well known fact that the effective aperture area of an antenna is directly proportional to its gain. Larger gain requires larger physical aperture, and this usually translates into a larger required spacecraft volume. With this interdependence in mind, several design compromises go into determining the type of antenna best suited for a particular spacecraft, always striving to minimize mass and volume. Innovative concepts are often required, but projects typically are unwilling or unable to allocate adequate resources for the associated technology development. In response to this funding challenge, the Telecommunications and Mission Operations Directorate Technology (TMOT) has been supporting three innovative spacecraft antenna technologies: inflatable reflectarrays, displaced-axis dual-reflector geometries, and layered-lens antennas (to be made inflatable, in the future). The ultimate goal of this effort is to develop space-qualified high-gain antennas with high efficiency, but reduced cost, stowage volume, and mass. The displaced-axis dual-reflector antenna has the potential to provide a compact configuration capable of yielding high efficiency with relatively small apertures (as small as 20 wavelengths in diameter). The inflatable reflectarrays and inflatable layered lenses address the need for very large apertures (300 or more wavelengths in diameter) with low mass, with application for both X-Band and Ka-Band high data rate telecommunication links.

Inflatable Reflectarray Antenna

Reflectarray antennas have been proposed as an alternative to reflectors [Berry et al. 1963], and have been studied at JPL, for spacecraft applications, for more than a decade [Huang 1990]. This technology consists of a (not necessarily flat) layer of scattering elements properly dimensioned to transform an impinging spherical wave into a reflected collimated beam. The particular reflectarray antenna of interest here was developed primarily to demonstrate that the required surface tolerance for Ka-Band operation is achievable with an inflatable 3-m diameter aperture. This diameter has been selected because it provides adequate link margin for future deep-space communications. A flat reflectarray is expected to maintain the required dimensional accuracy with less difficulty than a curved surface, such as a paraboloid, during long spaceflight missions. The antenna, shown in Figure 1, consists of three major components: a 3-m diameter, flat, thin-membrane reflectarray surface; an asymmetrical feed tripod; and a horseshoe-like torus that supports and tensions the 3-m thin-membrane aperture. Both the feed tripod and the torus are made of inflatable tubes with 0.25-mm thick, urethane-coated Kevlar material. The torus has a tube diameter of 25 cm, which was determined by the amount of tension needed to maintain the required membrane flatness (0.5-mm RMS). The 3-m reflectarray surface is a single-layer membrane having a 0.13-mm thick substrate made of polyimide material, with a 5-μ copper film deposited on both sides. On one side, the copper is etched to form approximately 200,000 patch elements. All of these elements are circularly polarized and identical, except for the rotation needed to achieve the required far-zone collimated radiation.

Test results show that the reflectarray antenna surface achieved a flatness deviation value of 0.1 mm RMS, which is significantly better than the required 0.5 mm. The measured radiation

Figure 1. Ka-Band inflatable reflectarray, 3-m

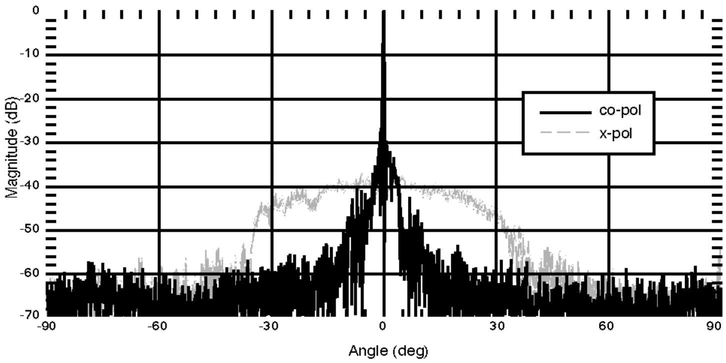

Figure 2. Measured radiation pattern of the 3-m Ka-Band inflatable array

pattern, shown in Figure 2, further substantiates the achieved flatness. The measured pattern has a co-polarized radiation with a beam width of 0.22 degrees, low average sidelobes (less than 30 dB below the co-polarization peak) and low average cross-polarization level (less than 50 dB below the co-polarization peak). The measured antenna gain of 54.4 dBi indicates an aperture efficiency of 30 percent, which is less than the expected 40 percent. However, several correctable areas have been identified to further improve the efficiency to 40 percent. The excellent radiation pattern and membrane flatness achieved indicates that inflatable reflectarray antennas at Ka-Band are feasible.

Displaced-Axis Dual-Reflector Antenna

Due to the high cost of launch vehicles, the current trend in space exploration missions is toward smaller spacecraft and, hence, smaller antennas. High-gain antennas using relatively small apertures (as small as 20 wavelengths) are usually implemented in practice by single-offset paraboloidal reflectors—in axially-symmetric paraboloidal geometries with small aperture sizes, the feed and associated transmission lines and supporting structures cause excessive blockage. Furthermore, conventional axially-symmetric dual-reflector geometries have severe subreflector diffraction losses caused by the small subreflector dimensions required to keep subreflector blockage at an acceptable level. Offset dual-reflector geometries are capable of good performance, but require more spacecraft volume. Axially symmetric dual-reflector antenna geometries can then offer advantages, if their limitations can be over-come. With this in mind, TMOT has supported a research effort on displaced-axis dual-reflector antennas. This basic antenna geometry, shown in Figure 3, appears to have been first proposed about four decades ago [Lee 1964]. However, its usage on small reflector-antenna applications is new.

The advantages of displaced-axis geometries, when compared to standard Cassegrain and Gregorian configurations, stem primarily from two facts: First, a central ray from the feed (aligned with the antenna axis) is scattered toward the edge of the main reflector by the subreflector; second, the main and subreflectors have a ring caustic (instead of a focus), located outside the rim of the horn. These two facts contrast sharply with what happens in Cassegrain and Gregorian configurations, where the central ray from the feed is reflected back into the feed, thereby causing high return loss, blockage, and undesirable diffraction characteristics. By design, the displaced-axis geometry minimizes the energy scattered back toward the feed by taking advantage of the natural feed radiation taper toward the subreflector rim. Also, since only a small amount of energy is scattered toward the feed, it can be placed very close to the subreflector, allowing for a large subtended angle, θ_E and, hence, the use of a very compact (low-gain) feed. The combination of these characteristics also allows the displaced-axis antenna to use small subreflectors (as small as four to five wavelengths in diameter) while still achieving relatively good diffraction behavior. Furthermore, a very compact configuration can be obtained by placing the backside of the subreflector in the same plane as the main reflector rim. The end result is a compact antenna with high efficiency, even when its dimensions would be too small for acceptable operation if used in conventional dual-reflector antennas. Figure 4 depicts a

CONTINUED ON NEXT PAGE

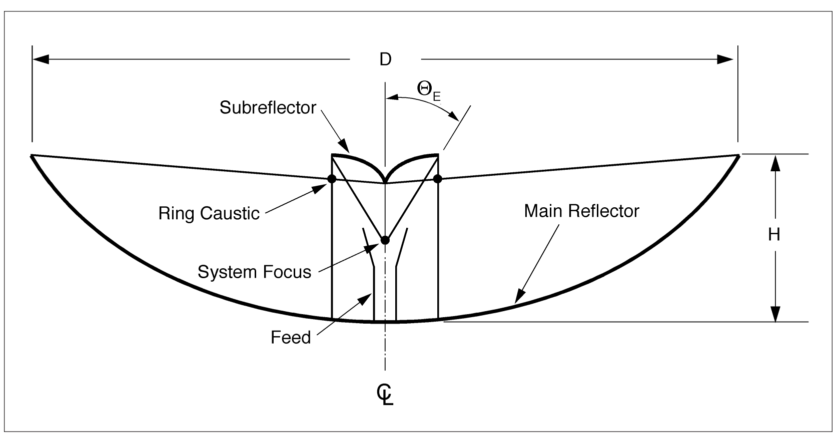

Figure 3. Displaced-axis dual-reflector configuration

CONTINUED FROM PAGE 9

prototype X-Band displaced-axis antenna. This antenna has a 1.25-m diameter aperture and yields a measured efficiency higher than 60 percent.

Because of the large subreflector subtended angle θ_E and the capability to operate with relatively small electrical dimensions, the displaced-axis antenna can accommodate both X- and Ka-Bands, provided that a suitable dual-band feed is used. To exploit this possibility, a special dual-frequency feed is currently being developed at JPL's Spacecraft Antenna Group. The feed uses two concentric circular waveguides, with the inner waveguide supporting the Ka-Band. The X-Band operates in the coaxial waveguide formed by the Ka-Band waveguide inner conductor and the outer circular waveguide. While compact and simple in concept, the feed presents complex design challenges, such as separating the X- and Ka-Band radiation, which spectrally are about two octaves apart. Furthermore, due to the reduced overall antenna dimensions, many of the commonly available electromagnetic modeling tools are not sufficiently accurate, making (numerically) rigorous analysis necessary.

The displaced-axis antenna effort has been divided in two different activities: reflector and feed developments. The reflector development uses rigorous electromagnetic diffraction analysis tools and benefits from the work on the initial prototype (Fig. 4). The feed development consists of designing a coaxial orthomode junction, a polarizer, and a dual-band coaxial horn. Each part of the feed is being designed and analyzed independently, but manufactured and tested as a system. Because of the uniqueness and novelty of the feed, several special components are also being designed and manufactured for conducting the

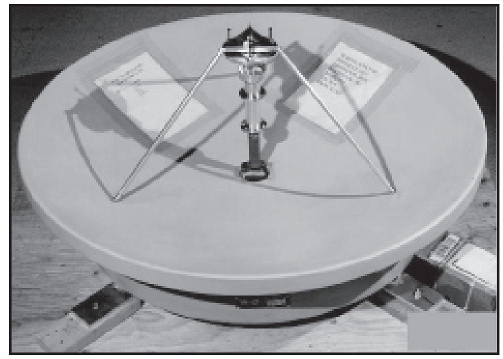

Figure 4. Breadboard displaced-axis antenna

appropriate testing required during the development stages. In order to minimize the number of breadboard units made and adequately predict feed performance, a combination of numerical and experimental tools are being used.

Layered-Lens Antenna

The implementation of large-aperture antennas relies on achieving high mechanical accuracy over the physical surface generating the aperture field. For a reflector to operate with tolerable gain degradation, typically the goal is (arguably) to achieve a surface accuracy better than 1/50 of the operation wavelength. This is achievable in principle, but it becomes increasingly difficult as the frequency of operation and/or aperture physical size increases. A promising method of implementing a high-efficiency aperture with reduced mechanical accuracy requirements is via a layered-lens antenna. In this geometry, conductive radiating elements are etched onto dielectric membranes (in this case, three), similar to those used in inflatable reflectarray antennas. The elements are dimensioned to transform the spherical wave produced by a radiating feed into a collimated

θ (deg)	Δ_ℓ (lens)	Δ_ρ (paraboloid)
0	0.00 δ	2.00 δ
30	0.15 δ	2.15 δ
45	0.41 δ	2.41 δ
60	1.00 δ	3.00 δ
75	2.86 δ	4.86 δ

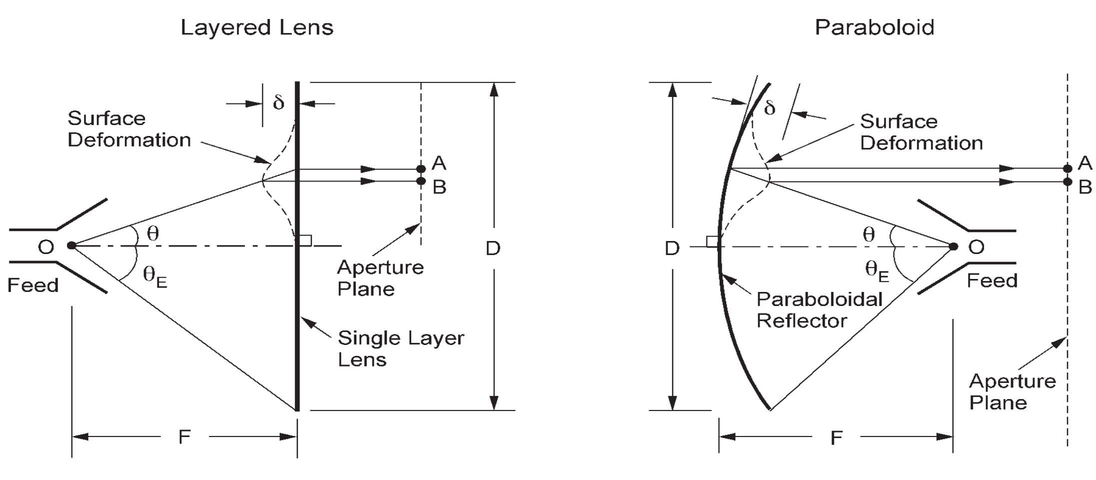

$$\Delta_\ell = \text{OA-OB} \cong \delta\left(\frac{1}{\cos\theta} - 1\right) \qquad \Delta_\rho = \text{OA-OB} \cong \delta\left(\frac{1}{\cos\theta} + 1\right)$$

Figure 5. Layered lens surface error sensitivity

beam as the energy propagates through the layered lens (Fig. 5). Note that the purpose of the layers is not to implement an artificial dielectric, which would require an excessive number of layers, but to effect the required phase transformation with the smallest possible number of layers. The reason layered lenses require reduced surface accuracy compared to reflectors can also be appreciated from Figure 5. In this illustration, the path length of a ray emanating from the feed and travelling to the aperture plane is shown for an axially symmetric layered lens and a paraboloid, with and without surface deformation. Basic geometry can be used to determine a first-order approximation for the aperture wavefront error of the paraboloid and the layered lens (shown in Fig. 5 as Δ_ℓ and Δ_ρ). Clearly, layered lenses are far less sensitive to surface errors than paraboloids. In fact, for the axial ray, a layered-lens surface deformation basically produces no wavefront error. Although surface deformations away from the axis of the layered lens produce progressively larger wavefront errors, this increase is offset by increased feed taper towards the lens rim, and by the fact that the lens support structure is fastened at the rim (smaller surface errors occur near the rim). Overall, it is anticipated that the layered lens has at least ten times less surface accuracy sensitivity than a reflector of the same dimensions—a Ka-Band lens would require less surface accuracy than a C-Band reflector. The layered-lens approach was first investigated almost two decades ago [Milne 1982]; however, its reduced surface accuracy property appears to have slipped by unnoticed.

A proof-of-concept Ka-band layered-lens antenna was designed, manufactured, and tested to spearhead a development effort on these antennas.

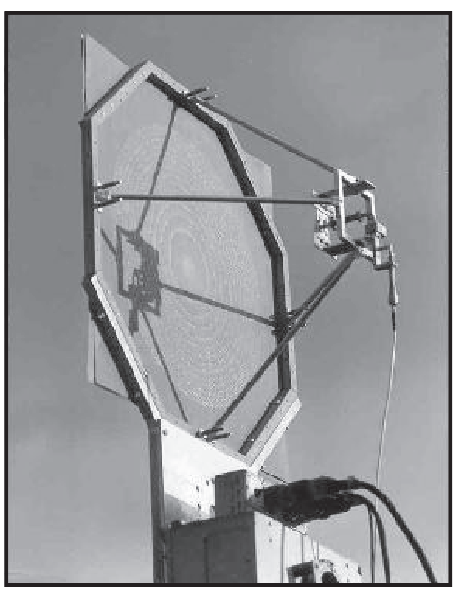

Figure 6. Ka-Band layered lens, 0.61 m

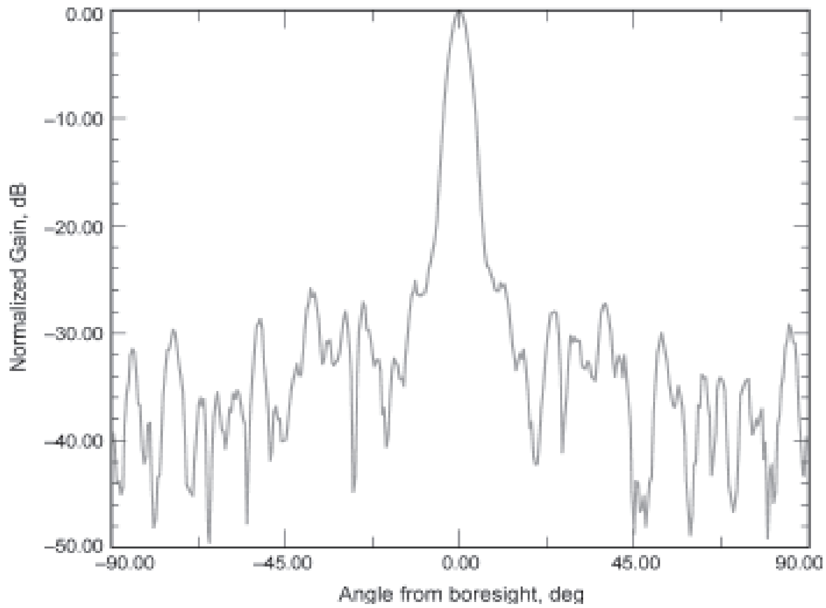

Figure 7. Radiation pattern of layered lens

The antenna has a 610-mm aperture diameter, a focal length of 305 mm, and uses three layers, with a total of approximately 140,000 top-loaded crossed dipoles to produce the required collimated radiation. Although layered lenses will ultimately be implemented using inflatable technology, the antenna shown in Figure 6 is semi-rigid. The measured efficiency of the layered lens was 34 percent and the associated radiation pattern results are shown in Figure 7. It was expected that this first lens would have low efficiency due to several approximations made in the analysis that lead to the top-loaded crossed-dipoles dimensions. Research effort is currently in progress to improve the analysis tools and, hence, bring the layered-lens antenna efficiency to values comparable to paraboloidal reflectors. It was not possible to verify, in a systematic matter, the insensitivity of the lens to surface deformations. However, preliminary evidence indicates that pushing the Ka-Band layered-lens center by as much as 10 mm impacts its gain by only a couple of tenths of a dB.

References

Berry, D. G., P. G. Malech, and W. A. Kennedy, "The Reflectarray Antenna," *IEEE Trans. Antennas Propag*., Vol. AP-11, 645–651, 1963.

Huang, J., "Microstrip Reflectarray Antenna for the SCANSCAT Radar Application," JPL Report No. 90-45, 1990.

Lee, J. L., "Improvements in or Relating to Microwave Aerials," U.K. Patent No. 973583, 1964.

Milne, R., "Dipole Array Lens Antenna," *IEEE Trans. Antennas Propag*., Vol. AP-30, 704–712, 1982.

Maximizing Science Return: Data Compression and Onboard Science Processing

Aaron Kiely, Roberto Manduchi, and Matthew Klimesh

Introduction

In the pursuit of deep space exploration we are constantly trying to increase the science data returned from each mission. As we strive to squeeze every dB out of the deep space communication link, it's important to keep in mind that maximizing the science return also requires both efficient data compression and intelligent data selection. That is, we must exploit onboard processing power to optimize both *how* data are transmitted, as well as *what* data are transmitted.

In the first part of this article, we'll get a glimpse of modern spacecraft data compression techniques. The second part of the article discusses how future missions will benefit even further when data compression is combined with intelligent onboard science processing methods. Our work thus integrates more conventional compression techniques with a science processing and buffer management component to yield a system that increases the science return of a mission.

Spacecraft Data Compression

It is clear that missions benefit when data compression techniques are used, but it may not be obvious *how much* of an improvement can be offered. Any gain obtained by employing data compression can be measured in dB—for example, data compression that provides a 2:1 compression ratio achieves the same benefit as a 3 dB increase in transmitter power. Figure 1 shows the savings in dB when the ICER image compression algorithm is used to compress a sample image.

With the recent emphasis on faster, better, cheaper missions, the use of commercial off-the-shelf (COTS) compression techniques might seem like a good idea. However, there are several reasons why the compression algorithm that works well on your web browser might not be so successful on a spacecraft:

- *Complexity*: Spacecraft processors are typically doing a lot more than just data compression, so low-complexity compression algorithms are essential. For typical Internet applications, the goal is fast decompression, while compression speed is not so important. On a spacecraft, the opposite is true.
- *Error Containment*: Without error containment techniques designed for the deep space link, a single bit error can corrupt large segments of data. Error containment schemes are not frequently used in common compression algorithms because terrestrial networks are able to use simple retransmission protocols to accommodate packet losses. Also, the packetization scheme used in deep space missions is different than that used on the Internet, and error containment strategies tailored to this structure can be expected to give better performance.

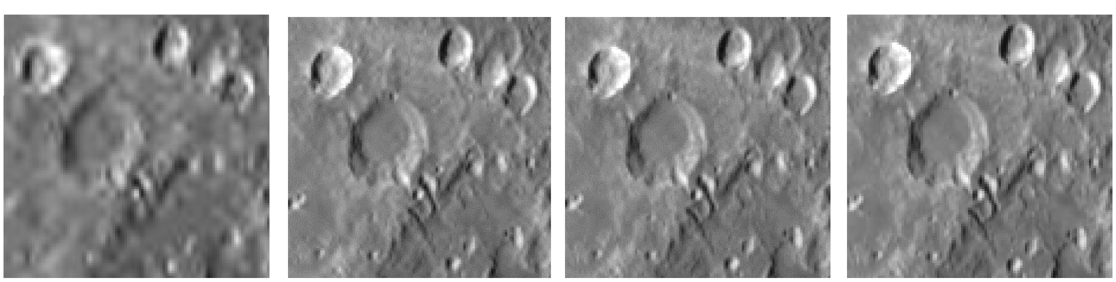

15.5 dB savings | **9.0 dB savings** | **6.0 dB savings** | **2.3 dB savings (lossless compression)**

Figure 1. This detail from a larger image shows an example of the effective link performance improvement offered by the ICER progressive wavelet image compression algorithm at various quality levels

• *User Requirements:* Generic data compression algorithms may not be suitable for scientific applications. The type of data, required quality, and computational resources available for spacecraft applications may be quite different from commercial scenarios. The compression algorithm that seems to do a good job compressing pictures of Aunt Bessie may not be so suitable for compressing pictures of the surface of Mars.

For example, the emerging JPEG-2000 image compression standard incorporates region-of-interest coding and accommodates low-memory implementations; however, it is not a low-complexity algorithm and it is not suitable for data types such as star field images or one-dimensional data.

One of the relatively recent trends in image compression is the use of progressive compression techniques. In a progressive algorithm, when any fraction of the compressed data has been received, we can reconstruct the entire image at nearly the best quality and resolution possible for that amount of data. Under the same circumstances, a non-progressive algorithm would only be able to reconstruct a fraction of the image. The ICER algorithm (a wavelet-based image compression algorithm developed in the Science Processing and Information Management work area) is an example of a progressive compression algorithm for deep space missions.

The use of progressive compression techniques allows for simpler and more effective use of downlink resources. Using non-progressive methods, users must adjust algorithm parameters to attempt to compress an image to the desired size. With a progressive algorithm a user can allocate space for an image and be assured that the resulting image quality is nearly optimum, subject to the allocation.

The ICER algorithm also incorporates an effective error containment scheme that exploits the progressive nature of the algorithm. For error containment, a compression algorithm must partition the data into segments and compress each segment independently of the others so that loss of data during the transmission of one segment does not affect the other segments. One method of doing this for image compression would be to treat an image as a collection of several smaller independent images that are compressed separately. Instead, the ICER algorithm incorporates a more sophisticated technique that does not segment image data until after wavelet transforms have been performed. The error containment technique used also trades rate between segments, so that regions of the image that are more difficult to compress are allocated more bits, resulting in improved compression and a uniform overall image. This also eliminates edge artifacts that would be seen under the more conventional approach—in the absence of channel errors, the error containment segmentation will be imperceptible to the viewer. The benefits of progressive transmission can also be seen under this error containment strategy. Figure 2 shows how an image might look after three segments have been affected by packet losses. Because the technique is progressive, even in the segments where much data is lost, low or medium quality versions of these segments can be reproduced.

A modified version of the ICER algorithm is also being produced that allows an onboard processing algorithm to specify regions of the image that should have higher priority, and therefore be produced with higher fidelity.

Science Processing and Buffer Management

While data compression plays a significant role in enabling deep space missions to maximize science data return over the constrained downlink channel, even the most sophisticated data compression methods become inadequate when the ability to collect sensor data on a spacecraft greatly exceeds the data transmission capability. Conversely, onboard data storage and processing capabilities are becoming increasingly affordable.

When a mission can collect much more data than can be transmitted to Earth, we must ask what part of the data should be downlinked, and what part should be discarded? We attempt to answer this question by:

- Defining a metric for the relative importance of gathered information; this is implemented by a suitable onboard science processing module;
- Making the best use of limited onboard memory resources via a prioritized buffer management mechanism which ensures that the more important data segments are transmitted to Earth with higher quality.

We can extend the idea of progressive transmission by incorporating semantic value to the "importance" attributes of encoded data. In particular, we define a simple measure of *science return* based on information-theoretic considerations as well as on the estimated *scientific value* of the transmitted data. The task of determining the relative scientific importance of segments of

CONTINUED ON NEXT PAGE

CONTINUED FROM PAGE 13

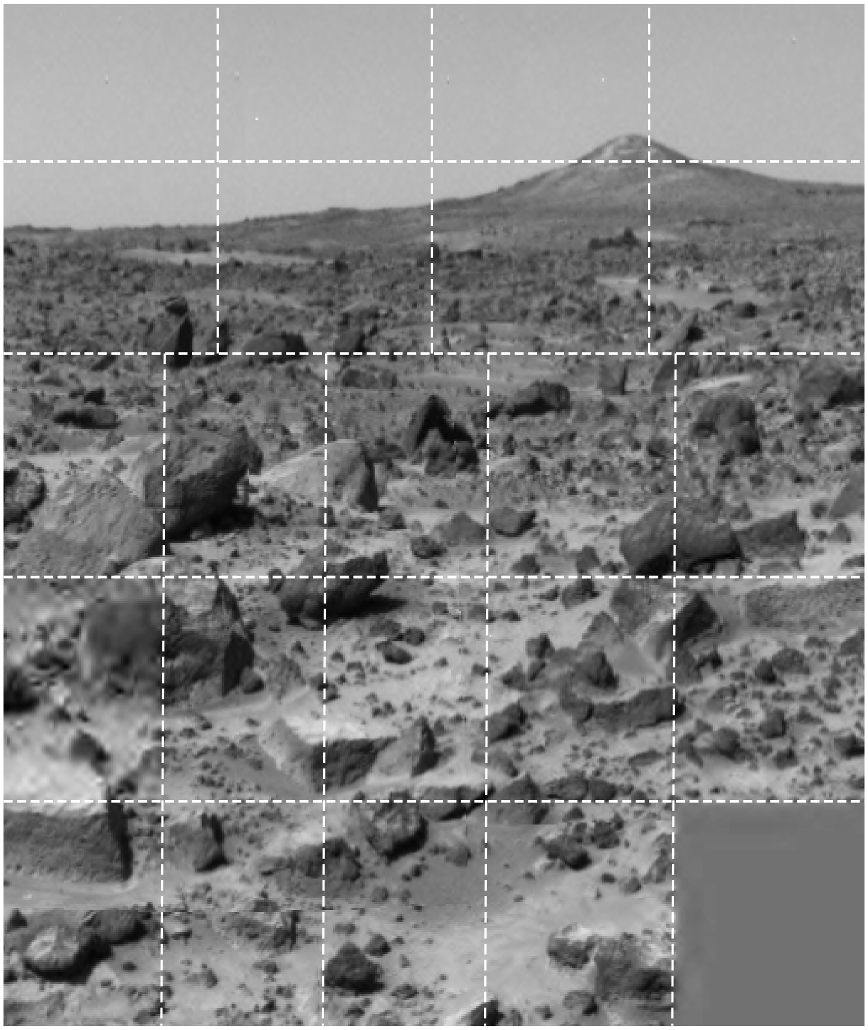

Figure 2: A 512 x 512 image segmented into 23 regions (shown by dashed lines) for error containment purposes and compressed to 2 bits/pixel. Three simulated packet losses produce varying amounts of loss in three regions of the image

data is carried out by an *onboard science processing* module, designed according to guidelines provided by the remote user (the science community). This module pre-processes the image and provides input to the progressive encoder in the form of a suitable classification map. By combining the semantic characterization produced by the science processing module with the content-blind data organization criteria of a traditional progressive encoder, we obtain a new measure of the importance of each segment of data. Algorithms that best utilize the available resources can thus be designed, and their performances assessed, in terms of science return.

The function of the science processing module is thus to assign priority values to the different parts of the acquired data. For example, in Figure 3 we show a sequence of images (first row) acquired at distinct times by a rover in a hypothetical experiment (the images were actually taken at the JPL Mars Yard). For this simulation, we assume a very simple set of science priority rules: "rocks" are more important than "soil"; and within the class "rocks" there are three level of priority: "basalt", "obsidian", and "shadow" (where there is very little textural or photometric information). It is understood that this is just one example of the possible prioritization rules that can be chosen—it will be up to the final user (the scientists) to decide on the most suitable science processing mechanisms and prioritization. In the second row of Figure 3 we show the results of image segmentation into the different selected classes operated by a statistical color-based classifier. The parts of the image labeled with black (basalt) have highest priority; those labeled with light gray (soil) have the least priority. The resulting image segmentation is used to drive a suitable region-of-interest (ROI) progressive compression algorithm such as our ROI extension of the ICER algorithm.

The priority information is used in our intelligent buffer management scheme. Our strategy is to allow the spacecraft/rover to take as many images as possible, and to transmit as much important information as allowed by the onboard and communication resources. In other words, we assume that the onboard buffer is constantly full and overflowing (any data packet which overflows is lost forever). Given the priority information produced by the science processing module, our buffer management scheme ensures that the high priority data packets in the buffer will be transmitted before the low priority ones, and that only the lowest priority data packets are discarded due to overflow. This strategy automatically adapts to the dynamics of the acquisition process, optimally exploiting the onboard resources. For example, in Figure 3 (third row) we show the received decoded data assuming that the downlink data rate was five times less than the acquisition data rate. The "important" image areas (rocks) are transmitted with much higher fidelity than the "uninteresting" parts (soil). In the fourth row we show the received data using a system *without* science prioritization. The amount of data received and the encoding mechanism are the same; in this case, however, a large number of bits is used to describe the "uninteresting" soil, and a much smaller number of bits are left to describe the "important" parts (rocks).

In our science processing scenario, scientists have less low-level control over the choice and quality of images returned, so it is important to verify that the automatic decisions can be made with sufficient accuracy. Related concerns apply when compression is used without science processing, unless the compression is lossless. With this idea in mind, a skeptical scientist once asked whether the serendipitous discovery of an erupting volcano on Io in Voyager images would have occurred if lossy compression had been in use. Compressed images showed that even at a ratio of 70:1, the volcano was clearly visible, which seemed like good news to the scientist—even at this high compression ratio, it was still possible to make significant discoveries. But, in fact, the news is even better—with a 70:1 compression we can transmit 70 times more images, maning scientists are 70 times *more* likely to make such discoveries!

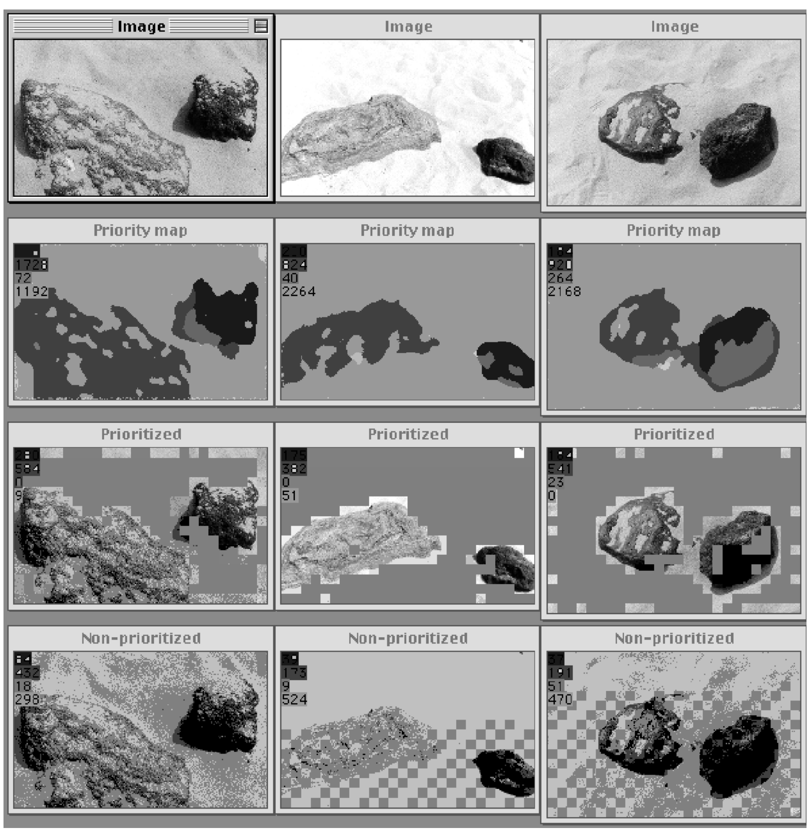

Figure 3: A simulation of science-prioritized data transmission: First row, original acquired data; second row, science-classified data; third row, data transmitted* with *our science-directed prioritization mechanism; fourth row, data transmitted* without *our science prioritization. The amount of data transmitted is the same in the prioritized and non-prioritized cases

CONTINUED FROM PAGE 2

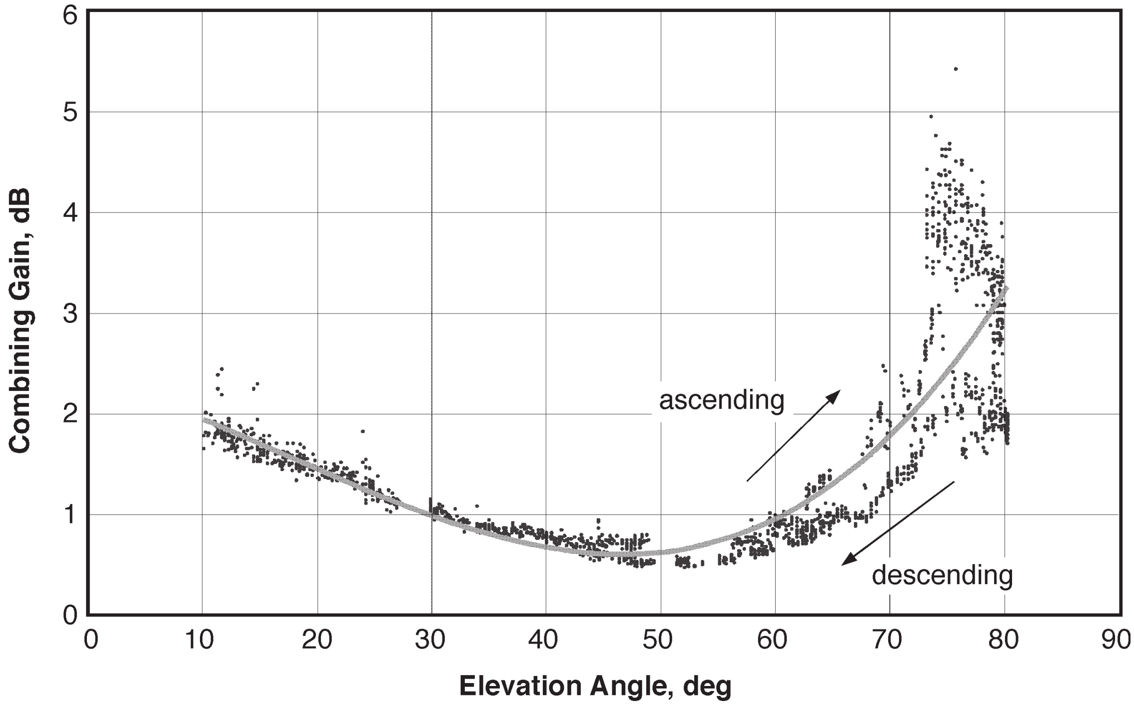

Figure 1. Measured performance of AFCS on 70-m DSS-14 antenna, 32 GHz

A representation of combining gain data recorded during the DS1 holography cone experiment is shown in Figure 1. Note that above 60 degrees elevation the data appears to split into two branches, indicating an azimuth dependence in the combining gain: the upper branch corresponds to the ascending part of the track (before transit) and the lower branch corresponds to its descent towards the horizon (after transit). This asymmetrical behavior of the antenna has been observed on all occasions, but the exact cause is not well understood at this time. The third-order trendline indicates average gain at a given elevation, without taking into account the azimuth dependence of the gain curve at higher elevations.

Array Feed Tracking with the Least-Squares Tracking Algorithm

The feed array can be viewed as an instrument that samples the focal-plane field distribution; therefore, it should be possible to extract pointing information from these samples in real time, even in the presence of antenna distortions. This hypothesis was tested during the holography cone experiments by measuring and recording the complex array response for a raster of small pointing offsets located around the optimum pointing direction at 10-degree elevation increments. Using the results of the raster scans, a coarse-quantized least-squares tracking algorithm was developed and installed in the DFPA to obtain preliminary information about AFCS tracking capabilities under actual operating conditions. While tracking DS1, closed-loop tracking capability and recovery from large applied offsets were both tested and verified on DOY-037 and 038, at elevations of approximately 46, 57, and 73 degrees. In each case, the least-squares tracking algorithm was able to accurately acquire and track DS1, and recovered successfully from manually applied pointing offsets of 8 mdeg or more on several occasions, as described in [Vilnrotter and Fort 1999]. Subsequent analysis [Arabshahi et al. 2000] of AFCS tracking performance using data obtained from DS1 tracks verified these experimental results. In addition, detailed analyses and simulations of AFCS least-squares and neural network tracking algorithms indicates excellent performance, even in the presence of severe antenna distortions and noise. For example, it was shown in [Mukai et al. 2000] using simulation that both least-squares and neural network tracking algorithms can achieve better than 0.5-mdeg rms pointing error at an SNR of 30-dB Hz, at both low and high elevation angles.

Deformable Flat Plate

DFP is a deformable mirror designed to correct gravity-induced antenna distortions on the main reflector surface by correcting wavefront phase errors. An experimental DFP was built with actuators laid out to enable gravity correction for either the 70-m or DSS-13 34-m DSN antennas [Bruno et al. 1996]. Its active surface was constructed from a 6000-series aluminum sheet, 0.04-inches thick and 27-inches in diameter, nominally producing a flat surface. It is driven by 16 actuators, via a look-up table, that change as a function of the antenna elevation angles. The open-loop calibration of the DFP was achieved by deriving the gravity deformation of the antenna at all elevation angles from holographic measurements taken at only three elevation angles: 47.2, 36.7, and 12.7 degrees [Rochblatt et al. 2000]. The computations assume that:

1. The antenna structural response due to gravity loading is linear.
2. The antenna response to gravity loading is symmetric relative to the antenna elevation-rigging angle.
3. The elevation-rigging angle (maximum gain) of the antenna can be accurately inferred from total power radiometry (TPR) efficiency measurements.

Note that in the absence of measurement noise, three angles is the minimum number of observations needed to solve for the three unknowns in the linear model used [Rochblatt et al. 2000].

In the usual case where noise is present, it is desirable to acquire measurements at more than three elevation angles so that a least-squares solution can be obtained. Conversely, if only three measurements can be made, making two of them near the extreme values of 0 and 90 degrees will produce the most accurate results. Due to the constraints of time and the highest viewable angle (47 degrees) available from this antenna location, neither of these conditions could be met at the time of the holography cone experiment at DSS-14; therefore, conditions 2 and 3 above were added to the derivation [Rochblatt 1999]. The values in the look-up table were computed by performing a ray trace from the deformed main reflector to the DFP position and the structural finite element model of the mirror [Imbriale and Hoppe 2000]. This process was repeated at 5-degree elevation increments to derive 17 sets of 16-actuator displacements. The effects of the applied DFP vectors were measured, iteratively corrected, and improved using the microwave antenna holography system at 12.7-degrees elevation [Rochblatt 1998]. Finally, the performance of the DFP was assessed by tracking natural radio sources and obtaining antenna temperature using TPR [Richter et al. 1999]. Positive compensation was achieved over all elevation angles, with average compensations at Ka-Band of 1.7 and 2.3 dB, at elevation angles of 10 and 80 degrees, respectively. These results successfully demonstrate an open-loop calibration of the DFP (Fig. 2).

Joint AFCS–DFP Performance Measurements

As we have shown, operating either the AFCS or the DFP singly provides some gravity compensation over all elevation angles. However, the performance of a combined AFCS–DFP system was significantly better than the performance of each one individually. Analysis of the data obtained on DOY-056 (Table 1), when DS1 was tracked continuously as it descended from an elevation of approximately 50 degrees to 8.5

CONTINUED ON NEXT PAGE

Table 1. Average combining gain of AFCS, DFP, and AFCS–DFP on DOY-056

Elevation, deg	G_{AFCS}, dB	G_{DFP}, dB	G_{joint}, dB	ΔAFCS, dB	ΔDFP, dB
8.5	2.1	1.8	2.9	1.1	0.8
13	1.8	1.6	2.5	0.9	0.7
16.5	1.6	1.4	2.2	0.8	0.6
23	1.2	1.1	1.8	0.7	0.7
28	1.1	0.9	1.5	0.6	0.4
33	0.75	0.6	1.0	0.4	0.25
38	0.5	0.3	0.6	0.3	0.1

CONTINUED FROM PAGE 17

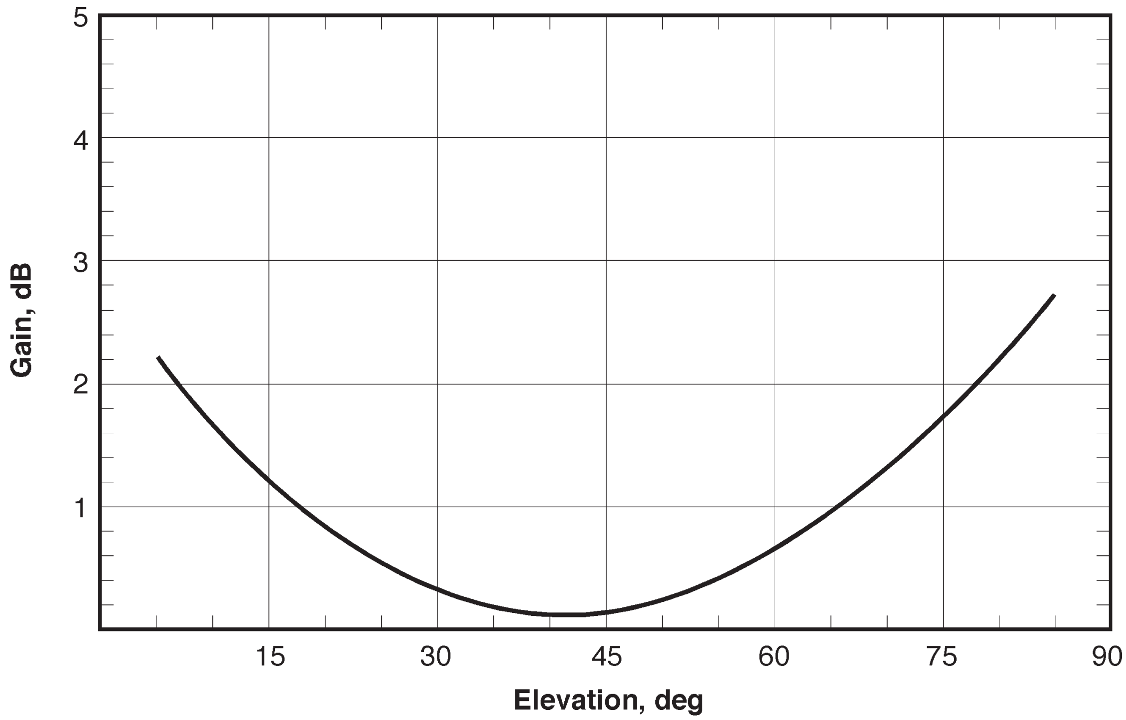

Figure 2. Measured performance of DFP on 70-m DSS-14 antenna, 32 GHz

degrees, clearly demonstrates that an AFCS–DFP compensation system working together in real-time can recover most of the signal energy lost to gravitational deformations on the 70-meter antenna at Ka-Band. The various gains used in Table 1 are defined as follows: $G_{AFCS} = G_c$ is the improvement over the central channel when the DFP is flat; G_{DFP} is the increase in signal power in the central channel when the DFP is activated; G_{joint} is the improvement over the uncorrected central channel when the DFP and the AFCS operate jointly to recover SNR losses; ΔAFCS is the contribution of the AFCS to G_{joint} over that of the DFP acting alone; ΔDFP is the contribution of the DFP to G_{joint} over that of the AFCS acting alone.

Summary and Conclusions

We have demonstrated that most of the SNR (or efficiency) losses incurred from mechanical antenna distortions can be recovered by means of a relatively simple real-time Ka-Band compensation system consisting of a seven-element array feed receiver operating jointly with a deformable flat plate. It has been shown that such a combination system recovers approximately 2.9 dB out of a possible 3.6 dB at low elevations, and is believed capable of recovering 5–6 dB (out of a possible 6.5 dB) at high elevations.

The present aperture efficiency performance of the DSS-14 antenna at Ka-Band is marginal as a result of two independent effects: loss at the rigging angle due to the rms surface error of the main reflector, and losses due to gravity-induced deformation (Fig. 3, lower curve, "Ka-Band Current"). The marginal rigging angle aperture efficiency of 31 percent is mostly due to the main reflector rms error of 0.73 mm. It is expected that this rms error can be reduced to 0.35–0.45 mm, thus increasing the overall antenna aperture efficiency to 55–60 percent at the rigging angle (Fig. 3, upper curve, "Ka Band + PNL SET"). The curve labeled "AFCS+DFP CURRENT" in Figure 3 shows the gravity performance of the antenna with joint AFCS-DFP compensation obtained from Table 1, while the curve "AFCS+ DFP+PNL SET" shows the expected efficiency performance when compensation is applied to the antenna after panel setting. Thus, with appropriate panel setting and gravity compensation benefits from both systems, the DSS-14 antenna can achieve excellent performance at Ka-Band, thereby enabling the predicted 6-dB communications improvement relative to X-Band.

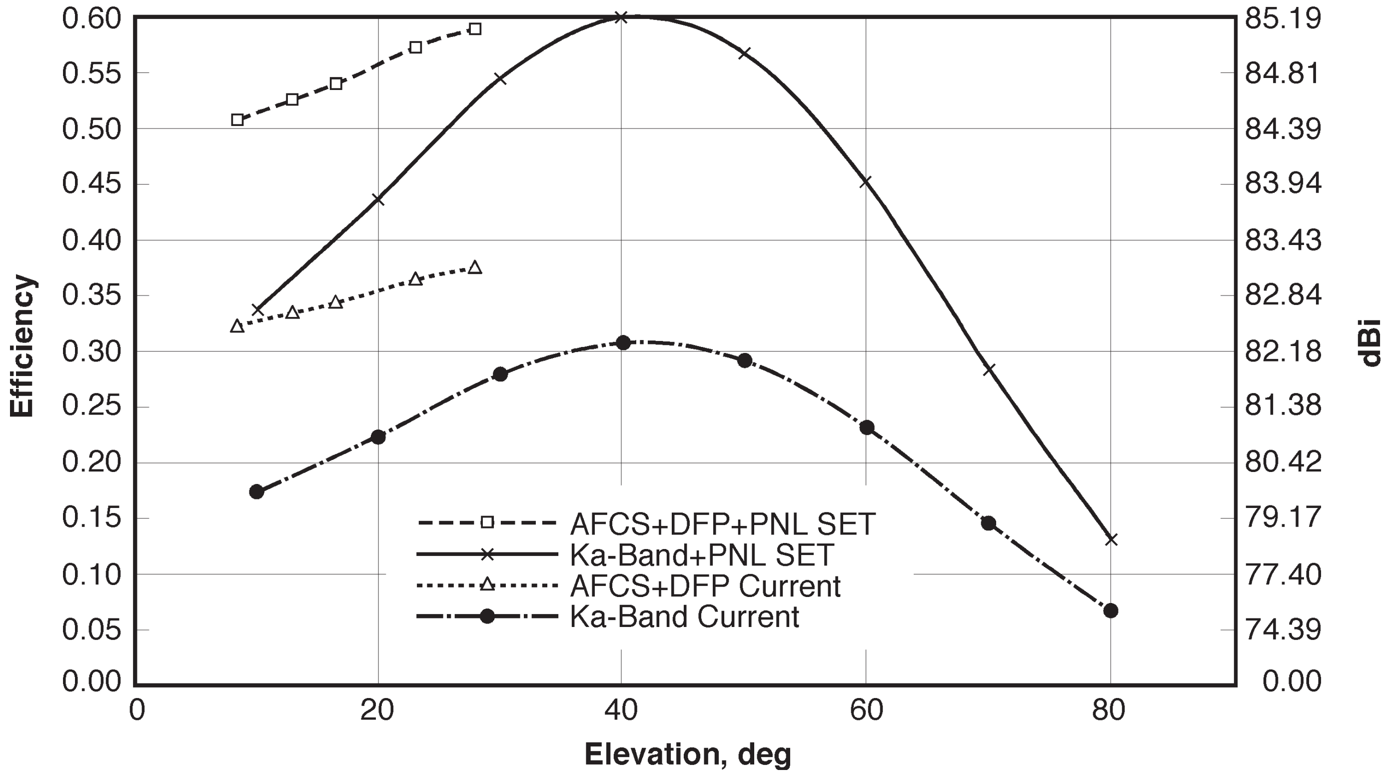

Figure 3. Performance of 70-m DSS-14 antenna, 32 GHz, present and future (predicted)

References

Arabshahi, P., R. Mukai, and V. Vilnrotter, "An Array Radial Basis Function Tracking System for NASA's Deep Space Network Antennas," accepted for publication in proceedings of *IEEE International Joint Conference on Neural Networks*, 2000.

Bruno, R., W. Imbriale, M. Moore, and S. Stewart, "Implementation of a gravity compensation mirror on a large aperture antenna," *AIAA Multidisciplinary Analysis and Optimization*, Bellevue, Washington, 1996.

Imbriale, W. A. and D. J. Hoppe, "Computational Methods and Theoretical Results for the Ka-Band Array Feed Compensation System-Deformable Flat Plate Experiment at DSS 14," TMO Progress Report 42-140, 2000.

Mukai, R., V. Vilnrotter, P. Arabshahi, and V. Jamnejad, "Computationally Intelligent Array Feed Tracking Algorithms for Large DSN Antennas," TMO Progress Report 42-141, 2000.

Richter, P., M. Franco, and D. Rochblatt, "Data Analysis and Results of the Ka-Band Array Feed Compensation System/Deformable Flat Plate Experiment at DSS-14," TMO Progress Report 42-139, 1999.

Rochblatt, D. J., "Holographic Measurements of the NASA-JPL Deep Space Network Antennas," Proceedings IEEE Aerospace Conference, Snowmass at Aspen, Colorado, 1998.

Rochblatt, D., "Ka-Band DFP/AFCS Demo on the 70-m," JPL Internal Report, http://knowledge2/descanso/, 1999.

Rochblatt, D. J., Daniel Hoppe, William Imbriale, Manuel Franco, Paul Richter, Phil Withington, and Herschel Jackson, "A Methodology for the Open Loop Calibration of a Deformable Flat Plate on a 70-meter Antenna," Proceedings of the Millennium Conference on Antennas & Propagation AP2000, Davos, Switzerland, 2000.

Vilnrotter, V., E. Rodemich, and S. Dolinar, Jr., "Real-time combining of residual carrier array signals using ML weight estimates," *IEEE Transactions on Communications*, Vol. 40, 604–615, 1992.

Vilnrotter, V., D. Fort, and B. Iijima, "Real-time array feed system demonstration at JPL," in *Multifeed Systems for Radio Telescopes*, Astronomical Society of the Pacific Conference Series, Vol. 75, 61–73, 1995.

Vilnrotter, V. and D. Fort, "Demonstration and Evaluation of the Ka-Band Array Feed Compensation System on the 70-Meter Antenna at DSS 14," TMO Progress Report 42-139, 1999.

ROVER OPERATIONS

Mars Rover Mission Distributed Operations

Paul Backes and Jeffrey Norris

Introduction

Internet-based rover mission operations will enable participation in planetary rover missions by an increased number of scientists and at reduced costs. The Web Interface for Telescience (WITS) and the Multi-mission Encrypted Communication System (MECS) have been developed to enable Internet-based Mars lander and rover mission operations. WITS provides downlink data visualization and sequence generation and MECS provides secure Internet-based communication. WITS and MECS were used by the Mars Polar Lander (MPL) mission for Robotic Arm sequence generation during operations readiness tests (ORTs) and during the post-mission field test. This article describes WITS and MECS and explains how they were utilized in the Field Integrated Design and Operations (FIDO) rover desert field test of May 7–18, 2000. Further information on WITS and MECS, including a public-outreach version that can be downloaded and run on a PC, can be found at the URL http://robotics.jpl.nasa.gov/tasks/wits/.

Internet-based operations provide various benefits. Use of the public Internet for distribution of mission data eliminates the costs of dedicated leased lines which have been previously used to support remote users. The MECS system makes it possible to use the Internet for data distribution by providing the necessary security and data encryption. Enabling users to participate in a mission from any location on the World Wide Web reduces both housing and operations center infrastructure costs. Expenses can be further reduced by enabling scientists to continue with other work at their home institutions during the mission. This is particularly important for long-duration missions and those continuing into extended mission phases. Internet-based operations are valuable for engineers and scientists at JPL; they can support multiple missions from one location, e.g., their office, and utilize computing resources outside of the mission operations area. Pre-mission ORTs can also benefit from Internet-based operations, because participants from other institutions can fully participate in ORTs without the additional cost (both in time and money) of having to travel to JPL for each test. Thus, more tests can be performed which involve all mission participants.

The FIDO rover is a prototype of the Athena rover that is being considered for the 2003 NASA Mars mission. FIDO is equipped with many of the elements of the Athena payload. It is used to simulate the complex surface operations expected of the Athena payload, focusing on identification of rock targets, approaching the targets and conducting in-situ measurements, and drilling and verification of cores. Stereo cameras on the rover's mast and body are used to image the surrounding terrain. An instrument arm places instruments on selected surface targets. A coring drill is used to extract rock samples.

The FIDO field test was used to evaluate new rover operations scenarios for the potential Athena mission, while utilizing Internet-based operations enabled by WITS and MECS. Rover operations are conducted independently of a lander and the rover can traverse long distances from the landing site. During the field test, operators used WITS and MECS to visualize downlink data and generate command sequences from JPL; Ithaca, New York; Birmingham, Alabama; Flagstaff, Arizona; St. Louis, Missouri; and Copenhagen, Denmark. In the operations readiness tests leading up to the field test, scientists Steve Squyres and Ray Arvidson and their colleagues from Cornell University and Washington University, respectively, were able to lead the tests from their remote locations. The field test demonstrated that WITS and MECS will enable efficient distributed operations in the next Mars rover mission.

System Architecture

WITS and MECS are designed to support multiple lander and rover missions. The architectures for their use in the MPL mission and FIDO rover desert field test are shown in Figures 1 and 2. For FIDO operations, WITS is used by itself. In the MPL mission, the APGEN (Activity Plan Generator) sequencing tool generated high-level

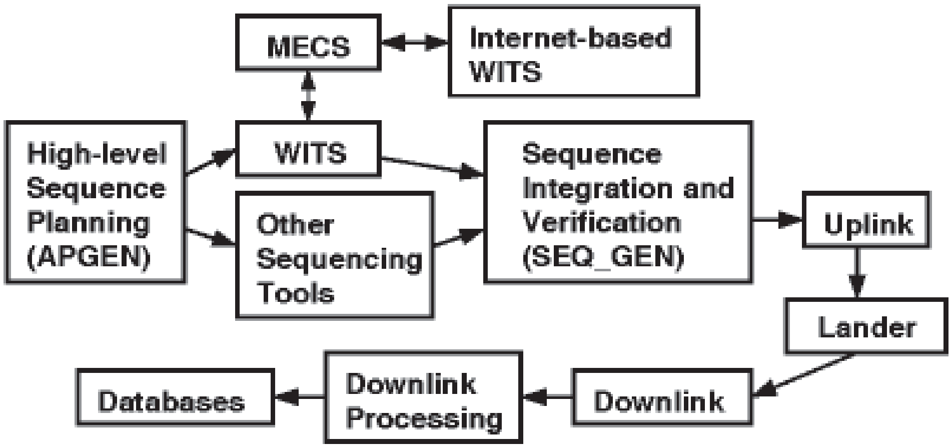

Figure 1. MPL Mission Operations Architecture

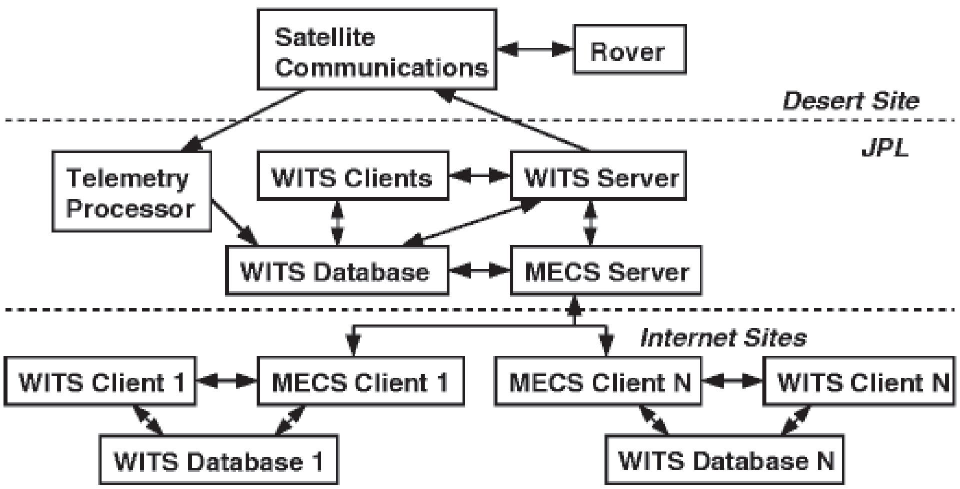

Figure 2. FIDO Operations Architecture

sequences. WITS added the low-level commands to satisfy the high-level goals within resource allocations for the Robotic Arm and Robotic Arm Camera. SEQ_GEN (Sequence Generator) was used to integrate all of the low-level sequences (including those for other instruments) into the final uplink sequence.

WITS includes a database, server, and multiple clients. The database is a structured file system that holds downlink data products and uplink sequence information. The server provides communication between clients and some database operations. Clients are distributed over the Internet and provide the interface needed to view downlink data and generate command sequences. Other tools, e.g., for planetary ephemeris or sequencing, can interact with WITS by reading and writing to the database or by direct communication with the server. WITS is implemented using the Java2 platform, including the Java3D and Java Cryptography extensions.

Internet Security

To enable collaboration in daily sequence generation by Internet-based scientists, a secure and efficient means to transfer data is needed. MECS was created to provide the required secure Internet-based communication. MECS was integrated with WITS for MPL and FIDO, and is designed to work with other mission tools as well. MECS operates in a fashion that is transparent to the remote user: Files simply appear as they

CONTINUED ON NEXT PAGE

CONTINUED FROM PAGE 21

become available and connections are made securely without any additional effort on the part of the user. MECS provides two types of encrypted communication for WITS: (1) automated delivery of database updates to Internet-based clients, and (2) encrypted communication between the clients and server. Since encryption was not required for the FIDO rover field test, the communication features of MECS were used without encryption.

MECS connections are authenticated using the NASA Public Key Infrastructure (PKI). After authentication, communications are made through SSL (Secure Sockets Layer) and are encrypted using the Triple-DES-EDE3 algorithm. Some of the advantages of MECS for Internet-based data distribution are: (1) MECS does not require remote users to request a data update; data is automatically delivered as it becomes available; (2) The MECS administrator can specify on a user-by-user basis exactly who will receive a particular type of file or directory; (3) Since MECS is written entirely in Java, it can run on most computers without modification; (4) MECS provides a high level of data security by using the SSL algorithm to perform authentication, and the Triple-DES-EDE3 algorithm for encryption; (5) MECS clients can be allowed to transfer files back to the mission control center. Files from users are stored on the server in a compressed, enveloped form that allows them to be scanned for hostile code. Since every client is authenticated using the NASA PKI, unauthorized users cannot transmit files to the server.

MECS is implemented as two Java programs, server and client, using the publicly available Entrust Java Toolkit's low-level authentication and encryption capabilities. For each mission, there is typically one server, operating behind a mission firewall, and many clients, one on each remote user's machine.

Figure 3 illustrates the steps necessary for a single MECS transaction. Steps 1 and 2 occur once, before the beginning of the mission, while steps 3 through 9 occur for every transaction. In step 1, a remote user must obtain a security profile from NASA PKI, which requires appearing in person at a NASA center security office. A security profile is a set of files on a floppy disk, protected by a password, that contain a user's private key. Users need their private keys to positively identify themselves online. The WITS server is also issued a security profile so that it can prove its identity to remote users. In step 2, each user must contact the mission administrator and

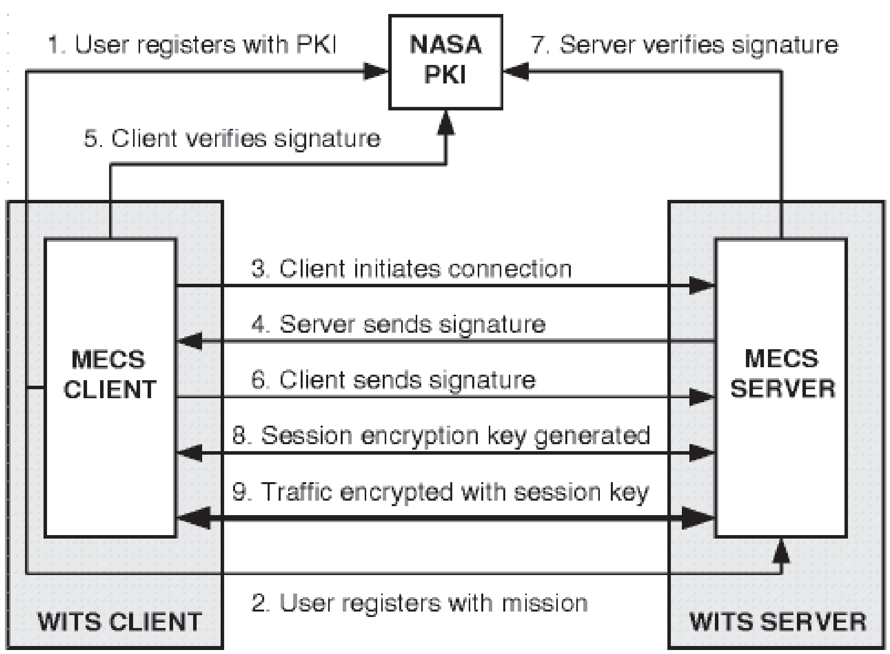

Figure 3. MECS secure communication steps

request that their profile be given access to mission data.

Steps 3 through 9 are repeated for every transmission from the client to the server or from the server to the client. In steps 4 through 7, the server and client exchange digital “signatures” generated from their security profile. They verify these signatures by communicating with the NASA PKI. This process, known as SSL authentication, relies on the fact that it is nearly impossible for someone to generate another user’s digital signature without that user’s private key and password. In step 8, the last step in establishing an SSL connection, the client and server use the Diffie-Hellman key agreement protocol to generate a unique symmetric encryption key that will be used to encrypt the traffic for this transaction. This encryption key is never sent as clear-text over the Internet, and all subsequent traffic for this transaction is encrypted, making MECS transactions invulnerable to eavesdropping or “packet-sniffing” attacks. In addition, every MECS transaction is protected by a new encryption key.

Downlink Data Visualization

Downlink data from the lander or rover is provided by WITS by means of various ‘views.’ Panorama, Overhead, Wedge, and 3D views with data from the FIDO rover desert field test are shown in Figure 4. The Results Tree window enables the user to select views to be opened based upon rover locations. The Plan window provides available views based upon a specific uplink plan. A plan includes view definitions and sequence information for generating one uplink sequence.

The Descent view provides images taken from orbit or during descent, and shows the landing location. The Overhead view shows the immediate area around the rover from above using various data formats (e.g., texture map, elevation map, contour map). The Panorama view is a mosaic of images taken by a stereo camera located on the deployable mast of the FIDO rover. The Wedge view displays one image of the panorama with various viewing options. The 3D view provides a

CONTINUED ON NEXT PAGE

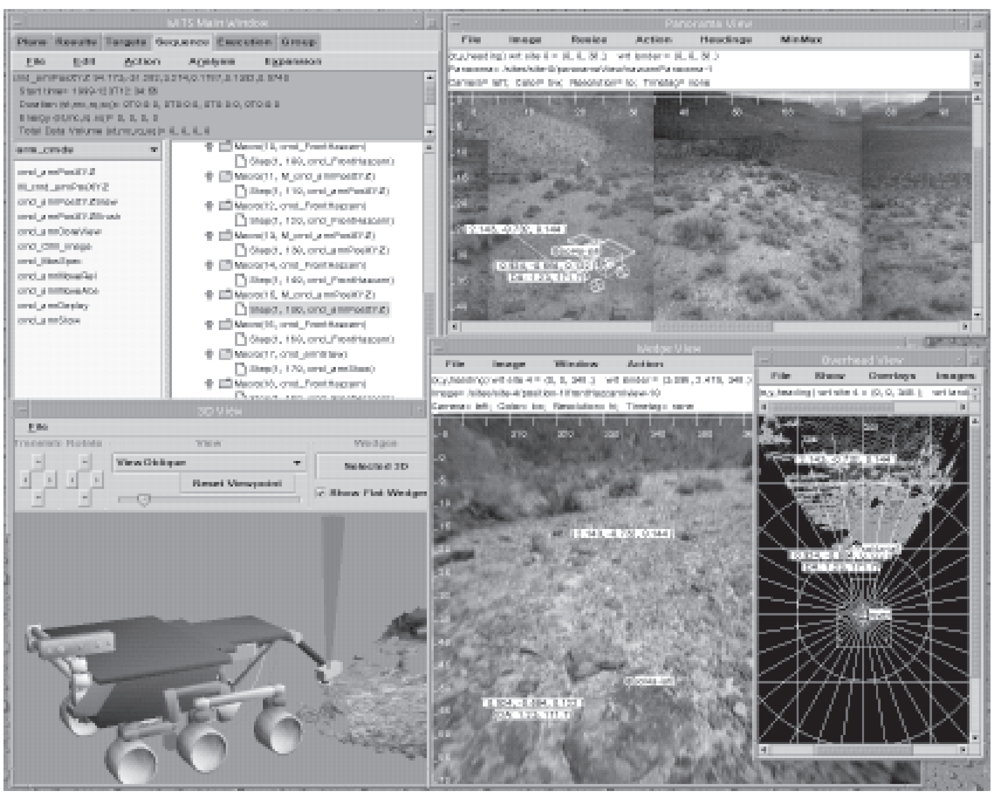

Figure 4. WITS visualization and sequence generation

CONTINUED FROM PAGE 23

3D solid-model visualization of the rover and terrain. Animated sequence simulation and state are visualized in the 3D view. The Contrast Adjuster view enables the pixel intensity range to be adjusted for a Wedge view image.

Selecting a point on an image causes the x,y,z position at that point on the terrain surface to be displayed. A target or rover waypoint can be created at the selected point via a menu option. Targets and waypoints are named 3D locations that can be used as parameters in sequence macros. Waypoints are generally used as locations for the rover to traverse through and targets are locations for science activities, such as drilling.

Sequence Generation

WITS provides various windows and features for command sequence generation. The Sequence window is used to generate a command sequence, which has a hierarchy of elements: Sequence, Waypoint, Request, Macro, and Step. A Request represents a high-level task. A Macro is the WITS element used to specify commands and parameters. Macros have expansions into Steps. A Step is a low-level command that will be uplinked to the spacecraft. WITS can generate output sequences in various formats, e.g., Spacecraft Activity Sequence File (SASF) format for the MPL mission and a FIDO-specific format for the FIDO rover.

The left side of the Sequence window displays a list of macros that can be inserted into a sequence. Multiple lists of macros are available. A Macro window is used to specify the parameters for a macro. A macro-specific algorithm expands the macro into one or more command steps.

A macro can generate View Objects that are displayed in the views to indicate what actions the macro is producing, e.g., footprints on the terrain for imaging commands. A valuable feature of the Sequence window is sequence state visualization. When a step in the sequence is selected, then the rover state is updated in the various views, with the state computed at the end of the step.

Resource analysis and rules checking are provided for sequence generation. The duration, energy, and data volume for each step of the sequence are computed and stored along with the cumulative duration, energy, and data volume at each step. Rules checking ensures that a sequence is valid relative to specified sequence rules.

Automated sequence report generation is provided. The sequence report is an HTML document that is produced with detailed information about a sequence. The report includes a description of the rover state at the start of the sequence and at each step of the sequence, screenshots of views at steps of the sequence, a resource report, the output sequence, and one movie for each view showing the rover motion through the sequence.

Distributed Collaboration

WITS provides features to support efficient collaboration within a distributed operations team. Users can input targets and command sequences and save them to the common server, as well as load the targets and sequences from the other users. A group mode is provided. If a user enters group mode, then the user can send and receive group commands. Group commands, which are executed on each group member's WITS client, include:

1. Open view, which opens the specified view;
2. Add marker, which displays a 3D marker with the originating user's name in all views; and
3. White-board commands that send messages to either one specific group member or the whole group.

An announcement feature pops up a window with a message in all group users clients. When a user joins the group, the state of the group, including all views and markers, is automatically initialized.

Acknowledgments

The work described in this paper was funded by the TMOD Technology Program, the Mars Exploration Technology Program, the SBIR Program, and the Mars Polar Lander Mission. The FIDO rover field test was led by Steve Squyres of Cornell University, Ray Arvidson of Washington University, and Paul Schenker, JPL task manager for FIDO rover development.

An Autonomous Diagnostic and Prognostic Monitoring System for NASA's Deep Space Network

Mark James, Ryan Mackey, and Lydia Dubon

Introduction

Our objective is to develop a system of integrated analytical and artificial intelligence methods that provide an autonomous and optimal solution for Fault Detection and Identification (FDI) decision-making in the DSN operational environment. Previous attempts have focused on neural-net-based systems that require large amounts of DSN training data and a priori knowledge of sensor data. Our system uses BEAM technology to not only process inputs currently used by human experts who perform diagnostics, but also to process all DSN data needed to achieve complete fault diagnostics and prognostics at every level of operations. In addition to our FDI-BEAM application, we have developed a SHINE-based 'expert system' that performs both model-based reasoning and case-based reasoning. The expert system has proved successful in exercising and validating heuristic knowledge using real-time DSN data.

To accomplish our goals we used two JPL-developed tools: Beacon-based Exception Analysis for Multi-missions (BEAM) and the Spacecraft Health Inference Engine (SHINE). BEAM is used as a highly advanced prognostic state estimator, and SHINE supports hard, real-time diagnostics and interpretation of the system state output by BEAM. These technologies provide new insights into system visibility that were not previously possible using channel-based diagnostics techniques, thereby making it possible to attain near zero false alarms. Raw sensor data and software-derived data are simultaneously fused in real time to automatically abstract system physics and information invariants (constants). This methodology enables a system to be ultrasensitive to degradation and changes, and to isolate significant events to specific sensors in both time and space.

The innovation of our approach is the synergy of BEAM's ability to provide complex DSN system analysis with SHINE's high-speed inferencing capability in a real-time continuous environment. This is accomplished using stochastic modeling, nonlinear information filtering, temporal channel analysis, and adaptive wavelet theory.

The goal of the FDI framework that we have built is to provide a DSN-compatible infrastructure that provides seamless integration of heterogeneous, intelligent tools for the purpose of DSN FDI analysis. The target domain for our prototype is the DSN antenna array configuration: It has been packaged as a DSN subsystem that interfaces with the DSN's new Network, Monitor and Control (NMC) computing environment. The FDI subsystem includes three components: (1) FDI server; (2) SHINE-based expert system; and (3) FDI-BEAM client application. The components are integrated into a tightly coupled architecture composed of a number of analysis components (Fig. 1). The FDI server software adheres to the DSN's standards and protocols for data exchange. The FDI server isolates tools such as SHINE and BEAM from the intricacies of DSN interface requirements. The SHINE-based expert system's rule base has been built incrementally with heuristic knowledge imparted by a DSN array expert. The SHINE inference engine is currently using CMU-LISP and dynamically links to the FDI communications wrapper. The FDI-BEAM client application consists of the core BEAM computational engine with the FDI wrapper. Both the FDI server and the BEAM client application have been prototyped using C. All diagnostic output is handled by a JAVA-based GUI.

Our FDI framework allows the expert system and BEAM client to continuously process DSN monitor data and predicts, and to generate diagnostics in a real-time environment. The automation that has been introduced is in the analysis and diagnosis, and is based on monitor data and predicted behavior. Currently, during spacecraft

CONTINUED ON NEXT PAGE

CONTINUED FROM PAGE 25

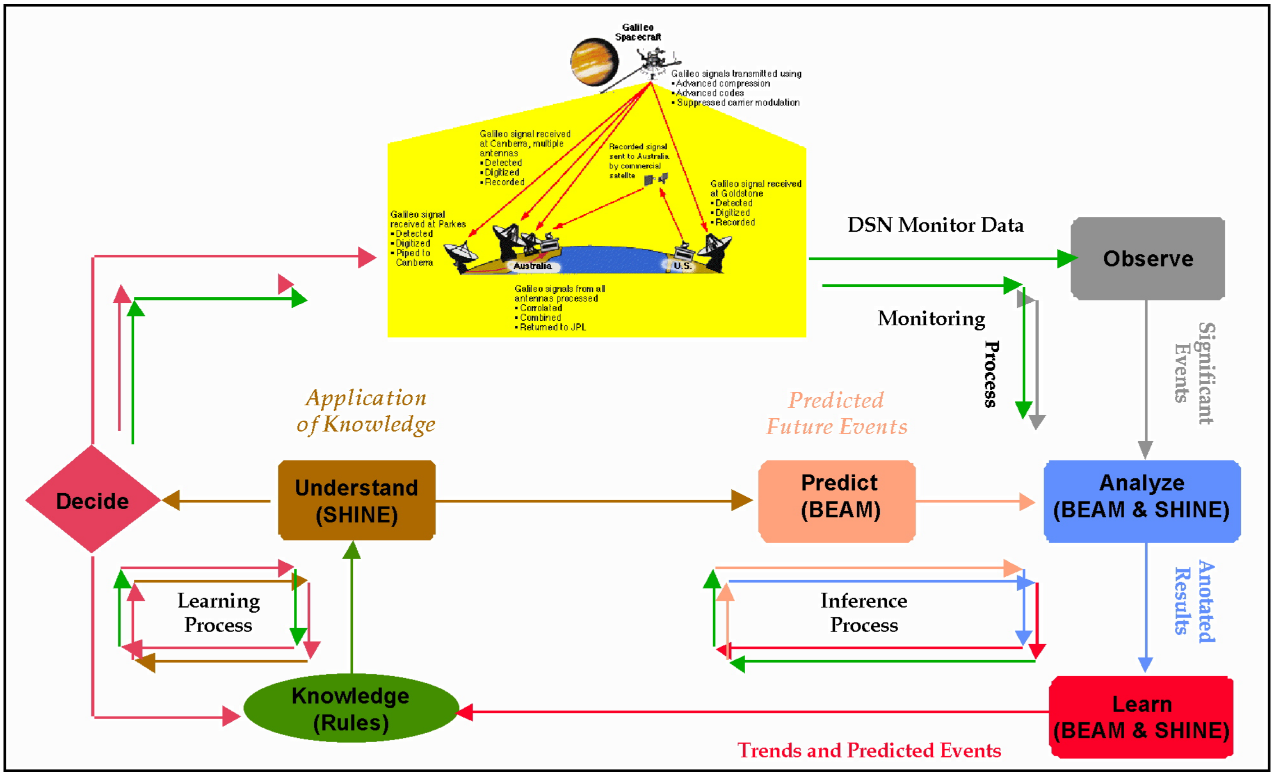

Figure 1. FDI system architecture

tracking operations, an operator must look at many displays of monitor data. From past experience, the operator focuses on changes in specific channels of data to determine if a problem exists. Our system has the ability to continuously process all monitor data, quantify monitor data contribution to system health, identify data signatures associated with nominal conditions, and detect and isolate anomalous events in an autonomous manner.

Discussion of BEAM

BEAM, which stands for Beacon-Based Exception Analysis for Multimissions, is a new method for automatic system analysis. It was originally funded by TMOD for use as an autonomous spacecraft monitor, but since then it has been adapted to a broad variety of applications. BEAM has a computational architecture and utilizes a series of reuseable software modules to perform real-time fault detection, system characterization, and impact analysis. The process is driven by data, though it can efficiently employ physical models, should they be available. Because there are so few other requirements, BEAM can be tailored rapidly to nearly any instrumented system. BEAM can be used as a key component of an end-to-end autonomous system; if desired it can also function in the lesser capacities of alert tool or analysis package. The FDI application here falls somewhere in the middle. The completed tool will serve alongside expert human operators, providing improved overall performance with reduced staffing. BEAM operates by studying regularly sampled performance data in real time. This data is studied singly, where individual signals are compared against their invariant properties, and in combination, where the sum total of sensor input is fused and system behavior as a whole is studied.

To characterize the overall system behavior, a complex transform is applied to multiple signals simultaneously, yielding a single, evolving object representative of the system physics. This object, which can be interpreted visually as well as mathematically, is tracked to reveal state transitions, emergence of faults and degradations, and acts on signals or recognized types of failures. Once these global properties have been extracted, further analysis is done on the implicated signals themselves in order to pinpoint sources of failure and progressions of incipient faults. The final results are tabulated and submitted to the interpretive elements of SHINE.

FDIR implementation of BEAM is a 3-step process. The first step is supervised training of the monitor. Known good data is fed through and the results stored in a highly compact abstraction. These results represent different modes of system operation, and can be studied to determine relative channel dependency and consistency, which is useful in system design. Alternately, if no system analysis is necessary, these automatically generated results can be fed immediately to the fault detector with no further operator interaction.

Once we have a workable training set, we are ready for runtime execution, and finally comparison and data extraction of the monitor's results. Training a fault detector is often a difficult ordeal, but we have managed to avoid most of the usual obstacles. In particular, once the basic configurational parameters of the system have been established, such as the number of signals, sample rate, etc., the monitor can be left to train "on-the-job." All training is done using system data, whether real or synthesized.

The output from this transform is called a Coherence Plot, which contains a great deal of information and is in itself useful to a trained operator. The transform also outputs a measure of System Stability. This is a single parameter extracted from the Coherence behavior over time representing the mode behavior of the total system. While these plots contain a wealth of information, they can be refined much further in order to produce an unambiguous and clear system assessment. Following the raw computation of system coherence, differences between the computed and expected coherence—encapsulated in the "Difference Plot"—are considered at mode transitions or upon significant deviation from nominal coherences (i.e., training results). The coherence difference is localized to signals that participate in the transition and, if unexpected or unrecognized, this represents the fault.

The results from separate signal analysis methods are combined and interpreted using the SHINE inference engine. Combination of these methods allows detection of complicated anomalies, the ability to distinguish between fault source signals and secondary effects, and permits a more sophisticated method of event recognition at the single-signal level.

Discussion of SHINE

SHINE, which stands for Spacecraft Health Inference Engine, is a multi-mission reusable knowledge-base software tool for the monitoring, analysis, and diagnosis of spacecraft and ground systems. SHINE advances the state-of-the art in artificial intelligence by enabling solutions to a broad class of problems that previously were considered intractable because of real-time system requirements, high-speed real-time and small-size constraints, portability or flight hardware ready availability. It introduces a novel paradigm for knowledge visualization and ultrafast forward and backward inferencing that goes well beyond traditional forward and backward chaining methodology.

SHINE provides more than a 1,000X increase in inference speed and up to 10,000X reduction in execution environment over the best comparable commercial product. SHINE's unique approach to knowledge representation and knowledge visualization and reasoning processes enables any developer to build high-performance expert systems, previously only within the reach of "AI experts." It is intended for those areas of inferencing where speed, portability, and reuse are of critical importance. Such areas would include spacecraft monitoring, control, and health; telecommunication analysis; medical analysis; financial and stock market analysis; fraud detection (e.g., banking and credit cards); robotics; and any other area where rapid and immediate response to high-speed and rapidly changing data is required.

SHINE is comprised of a collection of high-level software tools that revolutionizes the creation of efficient, reusable, knowledge-based software systems. Unlike existing expert systems, SHINE generates only application-critical code, thereby eliminating the need to deploy the complete environment. SHINE's unique approach to the creation of knowledge-based systems is based on sophisticated compiler technology, one of the fastest inference engines in the industry, and a large library of AI problem-solving techniques.

BEAM and SHINE In DSN Operations

The '99 FDI antenna array prototype successfully demonstrated: (1) Autonomous monitoring of downlink signal processing using DSN predicts and telemetry channel data; (2) Detection of anomalies based on SNR channel data analysis; and (3) Identification of DSN channel contribution to significant system behavior changes.

In our '99 work, we used BEAM to provide a framework of fault tolerance for key tracking observables, such as signal-to-noise ratios and range frequencies. BEAM was also used as a prognostic tool, indicating when the system was deviating from nominal performance requirements by detecting faults in parallel with the expert system. Our BEAM output consisted of: (1) A normage value, which is an instantaneous estimate

CONTINUED ON NEXT PAGE

CONTINUED FROM PAGE 27

of the system as a single-event-based metric of system health; (2) A normage limit, which is an instantaneous estimate of the system threshold; (3) A DSN channel ranked list, which identified the channels associated with significant system behavior; and (4) An operating map, which is a summarization of system channel behavior and provides a means of dynamic system state visualization.

The expert system performed further fault identification and isolation using heuristic knowledge. Forward chaining rules were used to define the semantics of FDI messages and to analyze telemetry frame data. Each data item is associated with one or more hypotheses, which are generated as data arrives. Backward chaining rules were used to resolve all ambiguities in the hypotheses generated during telemetry frame collection. Model-based reasoning was used to combine real-time channel data with conclusions generated by the backward reasoning phase, and to map conclusions to actual hardware configuration.

Figure 2 depicts simultaneous detection of anomalous events by the FDIR BEAM and SHINE components. Actual time-stamped Galileo tracking data was processed by both components in parallel. The anomaly in the data consisted of out-of-range symbol SNR values generated by the receiver subsystem. The block diagram displays a red receiver subsystem as a result of SHINE's detection of persistent symbol SNR out-of-range values. The Activity Plot depicts BEAM's detection of anomalous behavior as a normage value exceeding the threshold value also computed by BEAM. The associated BEAM ranking list contains residual symbol SNR and measured symbol SNR as the highest ranked channels contributing to the system's anomalous behavior. Thus, both technologies converged in their diagnosis of abnormal system behavior and asserted each other's conclusions.

Conclusions

Just as artificial intelligence can play an important role in the monitoring and diagnosis of mission operations, the tools that are used to develop such systems are also significant. This is especially true when issues of reliability, real-time performance, limited code and execution size, ease of use, and maintainability are factored in. The AI techniques and tools developed for the DSN are well-suited for the monitoring and diagnosis of ground systems in general. Both BEAM and SHINE run well in environments where system resources such as processor cycles and memory are at a premium. They have been demonstrated in stand-alone advisory systems for human operators, as well as components of embedded systems. Both of the tools generate C++ code that allows them to run efficiently in flight systems with real-time operating systems such as VxWorks.

The benefits afforded by the application of these tools and techniques are significant. The architecture and autonomous fault diagnosis techniques pioneered in the Spacecraft Health and Automated Reasoner Program (SHARP) system have demonstrated important benefits for operator productivity and spacecraft safety, and have the potential to reduce

CONTINUED ON PAGE 40

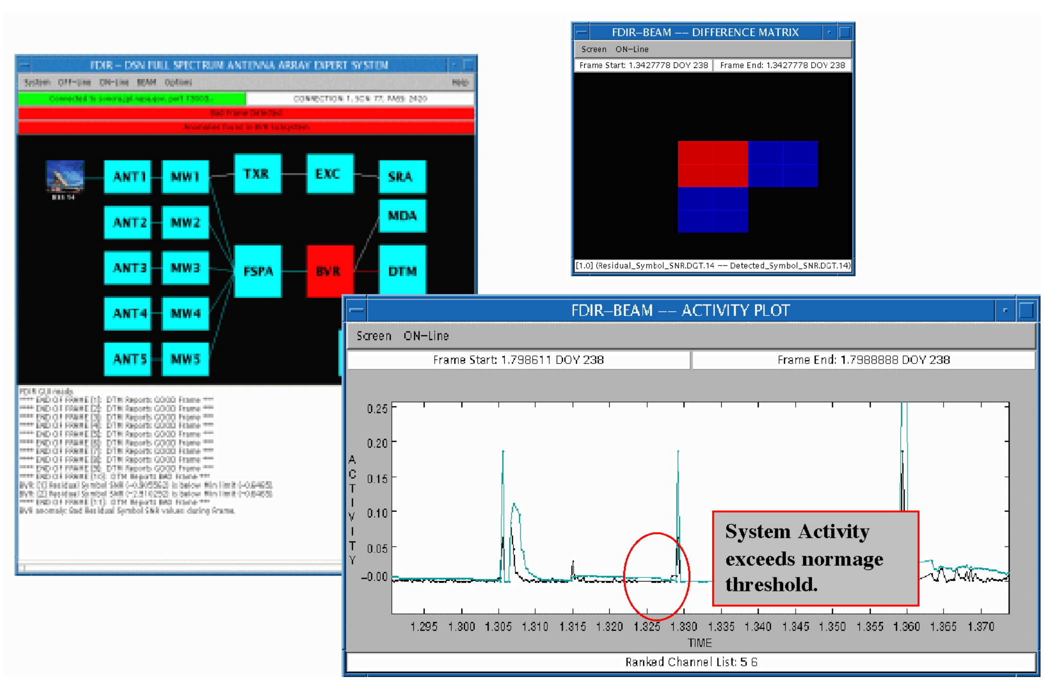

Figure 2. BEAM and SHINE during anomaly detection from Galileo pass 2420 archived SNRs

GAVRT Students Contribute to the Cassini Millennium Flyby of Jupiter

Mike Klein, Scott Bolton, and Michael Janssen

Introduction

December 30, 2000, marks the Cassini millennium flyby, when the Cassini spacecraft makes its closest approach to Jupiter as it swings past the giant planet on the way to Saturn. While instruments on the spacecraft conduct new observations of Jupiter from relatively close range, middle school and high school students from California and across the nation will be observing Jupiter using two ground-based 34-meter-diameter radio telescopes at Goldstone. The students and their teachers are participants in the Goldstone-Apple Valley Radio Telescope (GAVRT) science education project, which is a partnership involving the Lewis Center for Educational Research (LCER) in Apple Valley, California, the Telecommunications and Mission Operations Directorate (TMOD) at JPL, and the Apple Valley Unified School District [MOU 1999].

From October 2000 through February 2001, GAVRT students and their teachers will join an *ad hoc* research team of space physicists, radio astronomers, digital and microwave engineers, software engineers, mission operation specialists, educators, curriculum specialists, and numerous volunteers from many walks of life. The team is committed to deliver a prescribed set of radio astronomy measurements that will enable the Cassini spacecraft to perform previously unplanned observations.

Assembling the Team

How and why did such an eclectic group of people happen to join this science research team? Prior to launch, the scientists on the Cassini Radar Instrument Team decided to add a passive microwave radiometer to the radar receiver, which was on the spacecraft to perform radar imaging of Saturn's largest moon, Titan. The passive radiometer was built into the radar system by the Italian group that was responsible for the radar instrument. Michael Janssen, a member of the Cassini Radar Instrument Team, took responsibility for the calibration and science objectives of the radiometer. With the help of his colleagues, he determined that a carefully planned in-flight calibration scheme would enable the research team to achieve the best results from the radiometer observations of Titan and the rings and atmosphere of Saturn.

Titan is the only moon in the solar system with a significant atmosphere. In fact, Titan's atmosphere is deeper than Earth's, with a cloud cover so thick that its surface could not be imaged by the camera on the Voyager spacecraft. The Cassini radar experiment, operating at 13.8 GHz, is designed to penetrate the clouds and image Titan's surface, which is so cold that lakes of liquid hydrocarbons, such as ethane, are believed

CONTINUED ON NEXT PAGE

CONTINUED FROM PAGE 29

to populate it. The passive radiometer data will help identify the location, extent, and composition of these surface features.

The capability of the Cassini radiometer experiment prompted Scott Bolton, leader of the Cassini Jupiter Science Planning Team, to suggest using passive radiometry to map the high-frequency (and therefore, high-energy) microwave emission from Jupiter's radiation belts. This experiment will benefit from the flyby opportunity to "image" Jupiter's high-energy radiation belt close to the planet. The proximity of the spacecraft will facilitate adequate spatial resolution to map the synchrotron component of the microwave emission. Synchrotron radio emission is generated by extremely high-energy electrons trapped in the radiation belts. High spatial resolution maps are required to separate the synchrotron emission from the intense thermal emission of Jupiter's atmosphere, a task which has only been accomplished from the ground at much lower frequencies (<8 GHz).

Subsequent discussions with engineers and scientists from the Deep Space Network (DSN) Science Team led to a plan to use the radio astronomy flux-density scale to calibrate the spacecraft receiver. Nearly simultaneous observations of Jupiter's microwave emission at 13.8 GHz would be made from the spacecraft and from calibrated DSN antennas at Goldstone with Jupiter serving as the transfer radio source. The close approach of Cassini to Jupiter in December provides a unique opportunity to measure the intensity of Jupiter's microwave emission at levels that are well above the detection limit of the Cassini radar receiver.

With so many observations of Jupiter being planned, Michael Klein, TMOD Science advisor for the GAVRT project, suggested that the team be expanded to include the GAVRT project because remote observations of Jupiter are central to the Jupiter Quest curriculum. It was clear that GAVRT-trained teachers and their students would be "ready and willing" to participate. Furthermore, the GAVRT infrastructure has demonstrated the ability to support a campaign of remotely controlled radio astronomy observations of Jupiter and the selected calibration radio sources required by the project.

In January 2000 the team began work on the implementation plan. The "TMOD and GAVRT Support Plan for the Cassini-Jupiter Microwave Observing Campaign (Cassini-JMOC)" was prepared, reviewed, and then approved on May 1, 2000. With the approval, the team increased its effort to prepare for the observations, scheduled to begin in November 2000 and continue through February 2001.

The Role of GAVRT and the Lewis Center

The Cassini-JMOC observations at Goldstone will take full advantage of the remote observing capability developed for GAVRT. In collaboration with the TMOD DSN Science Office, the Lewis Center will provide experienced staff to coordinate, schedule, and remotely operate the GAVRT antenna at DSS-12 and the R&D antenna at DSS-13. Radio astronomers from the GAVRT project will provide advice, as well as support the planning and analysis of observations and the effective use of the antennas.

GAVRT teachers and students have a strong track record of producing high-quality observational results using DSS-12 from May 1997 to the present. The Cassini-JMOC plan is to select teachers who have already passed the GAVRT training classes and have run Jupiter Quest observing sessions with their students. This requirement will ensure that participating teachers are prepared and sufficiently experienced to provide a positive experience for their students and deliver quality data to the project. Student results will be reviewed by the GAVRT project before they are reported to the Cassini-JMOC science team, and ultimately, to the Cassini Program.

The educational outreach of Cassini-JMOC is a collaboration involving GAVRT and the Cassini Program. The goal is to provide opportunities for students across the nation to participate in the Cassini spacecraft flyby of Jupiter through classroom activities. Data gathered by the GAVRT students will be available on the World Wide Web. These data will form the basis for classroom lesson plans that can be implemented by teachers who have no previous knowledge of radio astronomy. Background materials will be provided and activities designed to enhance students' knowledge of Jupiter by introducing them to the concepts of radiation belts, polar auroras, and atmospheric thermal emissions. Lessons in comparative planetology showing the differences between Jupiter and other planets will be provided. Comparisons will be made between the atmospheres of Venus and the giant planets and Earth's atmosphere and cloud structure.

The goals of the Cassini-JMOC effort are to enhance the science return from the Cassini experiment and simultaneously provide new and unique opportunities for students to participate in an exciting scientific event.

Goldstone Observations

The ground-based observations in support of the Cassini-JMOC will be made at several frequencies spanning the microwave spectrum from ~1.4 GHz to 32.0 GHz (21 cm < λ < 0.9 cm). Four frequency bands will be observed at Goldstone and a proposal is being prepared to use the VLA to map the radiation belts at 1.4 GHz. The 34-m R&D antenna at the Goldstone Deep Space Communication Complex (DSCC) will be used for the critical observations at 13.8 GHz. The other frequencies observed at Goldstone will be 2.3 GHz, 8.5 GHz, and 32.0 GHz. Other radio astronomy observatories have expressed interest in participating in the observing campaign.

The Cassini-JMOC Engineering Team is designing a 13.8-GHz receiver system to be installed at the Goldstone 34-m R&D antenna, DSS-13. To minimize costs, designs for the RF and microwave-feed subsystems will be scaled from existing designs from the Holography Development Task. The RF amplifier will be an uncooled HEMT amplifier. The system will temporarily replace the X-Band receiver in the X-Ka position in the pedestal room. The existing X-Ka dichroic mirror will remain in place to provide simultaneous Ku-Ka receive capability at 13.8 GHz. The Goldstone multi-band observations will be used to:

- Separate the thermal and non-thermal components at 13.8 GHz, the frequency of the Cassini radar receiver.
- Monitor time variations of the synchrotron emission by continuing ongoing GAVRT Jupiter Quest and NASA/DSN Jupiter Patrol observations at 2.3 GHz and 8.5 GHz.
- Measure the absolute flux density of Jupiter at calibrated radio astronomy wavelengths and then interpolate the absolute flux density at 13.8 GHz.

Jupiter's microwave spectrum (Fig. 1) is comprised of two components. Thermal emission from the planet's deep atmosphere dominates the spectrum at frequencies greater than ~5 GHz,

CONTINUED ON NEXT PAGE

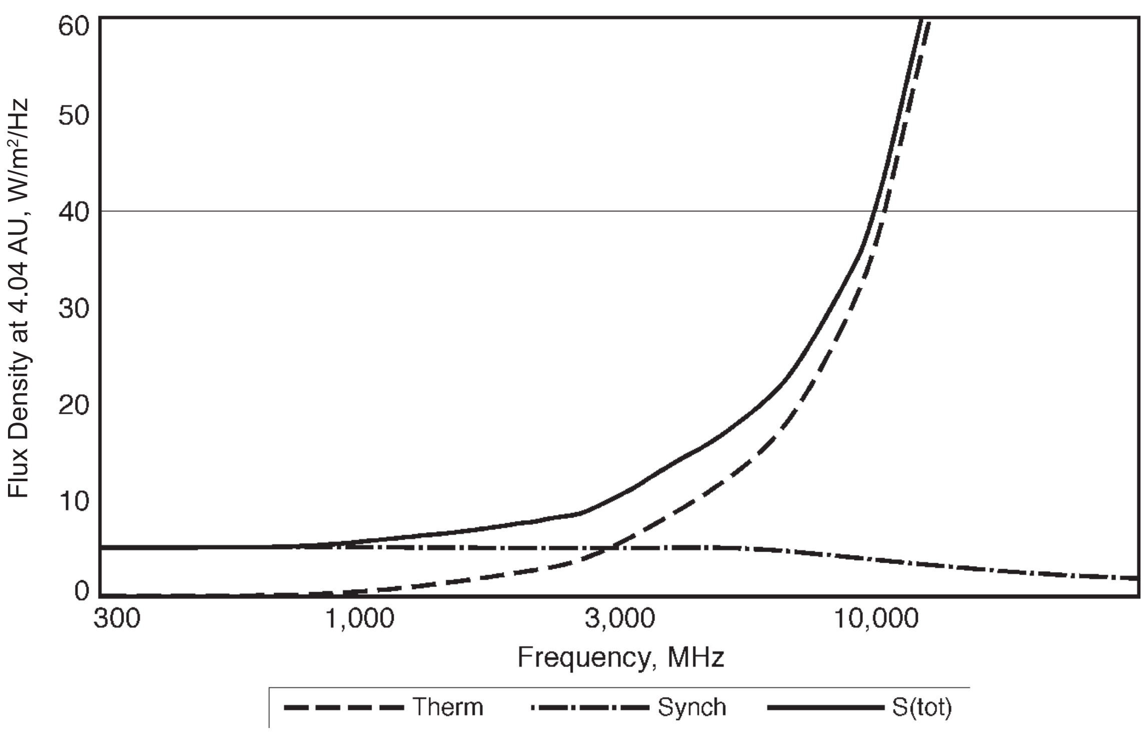

Figure 1. Jupiter's Microwave Spectrum (the sum of thermal emission from the atmosphere and non-thermal emission from the radiation belts)

CONTINUED FROM PAGE 31

while non-thermal (synchrotron) emission from Jupiter's radiation belts dominates the emission at lower frequencies. The non-thermal component is less than 10 percent of the total flux density at 8 GHz, which explains why ground-based microwave maps of Jovian synchrotron emission at frequencies above ~8 GHz are very difficult to achieve, and why none have been reported to date.

Cassini-JMOC Science Return

Cassini's encounter sequence at Jupiter has been designed to include using the Cassini microwave radiometer to observe and map synchrotron emission for two Jovian rotations a few days after closest approach on December 30, 2000. The high frequency (13.8 GHz) provides information on highly energetic electrons (20 MeV $< E <$ 50 MeV) that radiate at frequencies above ~8 GHz. These 13.8-GHz maps will be used to derive the spatial distribution of very-high energy electrons (>20 MeV) for the first time.

These "in-situ" observations from the Cassini spacecraft will be supported by ground-based Cassini-JMOC observations. Jupiter's synchrotron emission is known to be time-variable, and there is plausible evidence that the observed variations are correlated with changes in solar wind parameters, e.g., solar wind plasma density [Bolton et al. 1989]. The Cassini encounter with Jupiter will occur as the Sun's activity is reaching its peak in the current 11-year solar cycle. The last large-scale (>20 %) increase in Jupiter's synchrotron flux density, not counting the 1994 impacts of the Shoemaker-Levy 9 comet, was observed in the winter of 1989-90, just about eleven years ago. If the correlation holds, one might expect the synchrotron component to surge in the near future.

Cassini-JMOC ground-based observations will include frequencies below 8 GHz to monitor the non-thermal synchrotron component of Jupiter's microwave radio emission. The results of these observations may prove to be especially valuable if the intensity of Jovian synchrotron emission changes during the encounter. Knowledge of time variability will also be important to achieve the precision required for the receiver calibration.

Furthermore, it would be useful to have multi-band observations of time-variable synchrotron emission to compare with the results of new solar wind observations. The Cassini and Galileo spacecraft will both be searching for solar wind correlation with Jovian magnetospheric and atmospheric phenomena for approximately six months (from October 2000 to March 2001). These observations are expected to produce an unprecedented set of solar wind data in the vicinity of Jupiter; the study of synchrotron variability with the Cassini-JMOC could make an important contribution to the study.

The Cassini-JMOC team is anticipating some exciting times next fall and winter when the students are "on line" taking data during the Cassini Millennium Flyby of Jupiter. Having GAVRT students across the nation conduct radio astronomy observations that will be used to enhance the science return from the Cassini spacecraft is good news for the students, for NASA, and for public education.

References

Memorandum of Understanding, JPL MOU-99-720-149, 1999.

Bolton, S. J., S. Gulkis, M. J. Klein, I. DePater, and T. J. Thompson, "Correlation Studies Between Solar Wind Parameters and the Decimetric Radio Emission from Jupiter," *J. Geophys. Res.*, 94, 121–128, 1989.

****

For information about the Lewis Center or GAVRT, visit the websites at:
http://deepspace.jpl.nasa.gov/dsn/applevalley
http://www.avstc.org

Toward Radio Detection of Ultrahigh-Energy Neutrinos Using NASA's Deep Space Network

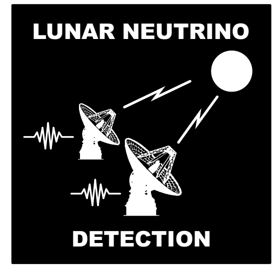

Peter Gorham, Kurt Liewer, Charles Naudet, and David Saltzberg (UCLA)

Scientific Motivation

Any vision for the future of both space- and ground-based astrophysical research must include efforts to establish a neutrino-based astronomy. This is the goal of an exploratory JPL–UCLA experiment now in its initial stages of operation using Goldstone Deep Space Network (DSN) antennas.

Virtually all of our information about the universe, the solar system, and even Earth on a global scale comes to us now as electromagnetic radiation, via the particles we know as photons. However, in terms of the sheer number of cosmic particles incident on Earth, neutrinos far outrank any other form of incident radiation. In fact, there are at least three species (or "flavors," as they are known) of neutrinos, which, along with their antiparticles, form six possible distinct types of radiation. Each of these neutrino forms can arise from different astrophysical processes, and signal a different physical environment in the source. Neutrinos also share one important property with photons: they travel in essentially straight lines from their source to the point of detection. This property allows us to form images of the source structure, thus attaining the requirements of true astronomy, if we can learn to build efficient detectors for these elusive particles.

We can be certain that there are copious sources of neutrinos in almost every class of astrophysical object. Two such sources, the Sun and the supernova 1987A in the Large Magellanic Cloud, have been detected. The neutrinos from these two sources are relatively low-energy particles (1–30 MeV) and are thus extremely difficult to detect from any other types of sources. If, however, one can extend the sensitivity to the high-energy neutrino domain, there are likely to be many more sources visible: pulsars and supernova remnants, active galactic nuclei, gamma-ray burst sources, and even exotic relics of early universe inflation. All of these are expected to be strong neutrino sources at energies above 10–100 GeV.

Of even greater potential significance are predictions of a diffuse ultrahigh-energy (UHE) neutrino background. These diffuse neutrinos may in fact hold the key to understanding the source of UHE cosmic rays, which are charged particles (protons and nuclei) that constantly rain down on Earth from space, and whose sources are still a complete mystery. Detection of such UHE neutrinos may even signal the presence of distortions in space-time, known as *topological defects* that survive from the very first moments of the Big Bang.

It is now widely believed that neutrinos with energies in the UHE range (10^9 GeV—a billion billion times the energy of visible-light photons) may, under plausible circumstances, provide the most accessible detection window for astrophysical neutrino sources other than nearby supernovae. This is due to the fact that the background atmospheric neutrinos, which dominate measurements in the 1000-GeV range and below, are no longer important at higher energies. Of equal significance is the fact that much of the secondary pulsed emission from the particle cascade that results from the neutrino interaction is largely coherent (thus growing quadratically in power with particle energy) in the radio regime. Above 1 million GeV the number of radio quanta exceeds the number of optical photons emitted, and at higher energies the radio emission is expected to dominate all other forms of secondary radiation products. This latter effect provides strong motivation to develop time-domain radio detection techniques for such cascade pulses; this effort is a major component of our experiment.

CONTINUED ON NEXT PAGE

CONTINUED FROM PAGE 33

Coherent Radio Emission from High Energy Cascades

The idea of detecting high-energy particles by means of the nanosecond pulse of coherent radio emissions from the cascade they produce can be traced back nearly 40 years to G. Askaryan (1962, JETP 14, 441; 1965 JETP 21, 658). Askaryan even suggested that this property could lead to neutrino detection by utilizing an array of antennas that sensed emission from upcoming events within the radio-transparent outer 5–10 m of the lunar regolith. Although significant early efforts were successful in detecting radio emission from high-energy particle cascades in Earth's atmosphere, surprisingly little work was done on Askaryan's suggestions that solids, such as ice and the lunar regolith, could be equally important media for detection. However, in the mid-1980s two Russian investigators, M. Markov and I. Zheleznykh, revisited these ideas and confirmed the theoretical basis. In 1988 Zheleznykh made a bold suggestion—that the large antennas used in ground-based radio astronomy should in fact be able to detect the radio pulses from UHE neutrinos interacting in the Moon's surface layers. Further evaluations of this suggestion have confirmed the theoretical basis.

Of course, such events should happen on Earth as well, but the problem is that at such close range it is difficult to "see the forest because of the trees." The main reason for current interest in this lunar-observation approach is that the predicted radio emission, which peaks in the microwave region, is well-matched to typical capabilities of the large antennas in the DSN and other large ground-based installations. In turn, the beam sizes of these antennas are of the same order of magnitude as the Moon's angular size in the sky. This means that the antenna is able to simultaneously monitor an enormous volume of material—roughly the visible area of the Moon times the approximately 10-m depth of its surface regolith layer. Even after accounting for inefficiency in the microwave transmission and detection process, the effective volume for neutrino detection is equivalent to about 100,000 km^3 of water—far greater than any ground-based detector can now monitor by looking at local material. Although systems that monitor terrestrial materials (typically large bodies of water or ice) are effective at much lower neutrino energies, they can never hope to achieve sensitivity at ultrahigh energies where enormous volumes are required.

Accelerator measurements

An important component in developing radio detection techniques for high-energy particles is to validate our understanding of the cascade radio emission process, which is to date largely untested experimentally. Particles of energy greater than several thousand GeV cannot presently be produced in accelerator facilities, but there are alternative ways to test the radio emission mechanisms using lower energy, pulsed, charge-particle bunches at a linear accelerator. We have performed a proof-of-concept experiment along these lines.

The experiment was performed at the Argonne Wakefield Accelerator in Argonne, Illinois, during September of 1999, using 20-ps to 50-ps pulsed electron bunches with charges of 1–20 nC. Detection of the radio emission was done using a pyramidal horn tuned for a frequency range of 1.7–2.6 GHz, as well as a balanced dipole with similar frequency range. The resulting radio pulses were detected and digitized using a real-time digital oscilloscope with 10 Gs/s (3-GHz bandwidth). Thus, we were able to directly sample the resulting electric field strength of the pulses, which produced coherent radio emission as they were dumped into a silica sand target designed specifically for the experiment.

In Figure 1, we show some of the resulting pulse profiles measured in this experiment. The voltages have been converted to the electric field strength and referenced to 1-m distance, making use of the known effective height of the antennas. The pulses show extremely high field strengths, in the range of 1000 V/m at these distances, consistent with predictions based on Askaryan's hypothesis. Thus, although these measurements cannot yet be traced back directly to a UHE cascade, they do lend strong support to the theory of how the coherent radio emission develops once the cascade has formed.

Neutrino interactions in the Moon

At energies of 10^{18} to 10^{20} eV, the mean free path to interaction for a neutrino in the lunar crust varies from 500 km down to 60 km. Because the radius of the Moon is 1740 km, detectable neutrino interactions at these energies will involve particles that traverse moderately to short chords of the lunar crust. Since there is equal probability of interaction per unit pathlength, we then detect the fraction of neutrinos that interact near the end of the chord just below the lunar surface. All of

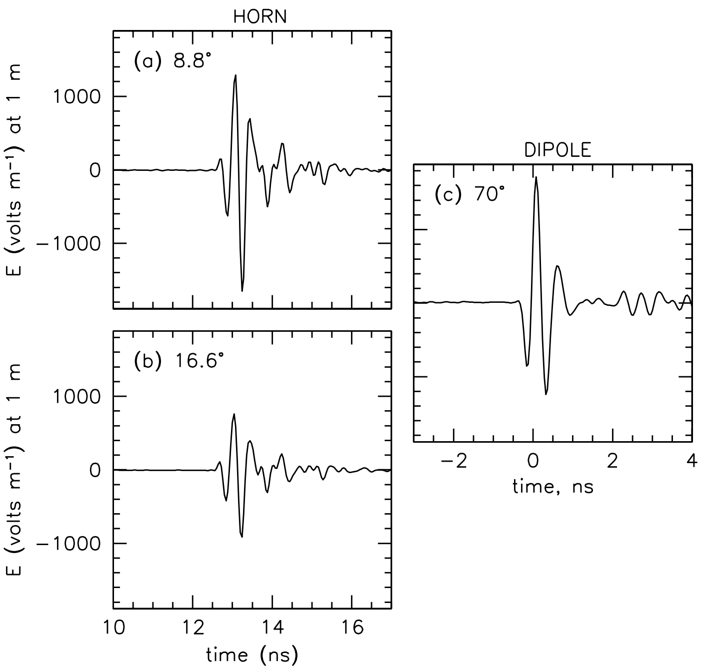

Figure 1. Coherent radio pulses from accelerator measurements at Argonne National Lab. Note the extremely high electric field strengths measured. The time differences in the plots are due to uncorrected cable delays; the pulses are actually precisely coincident in time

the neutrino flavors at these energies will primarily produce a particle cascade that yields coherent radio Cherenkov emission.

Figure 2 shows the geometry for coherent radio Cherenkov production. Since the neutrino track angle is upcoming relative to the tangent to the local surface, the forward-beamed radio emission can escape through the regolith interface, after significant refraction due the change of the refractive index from the regolith to vacuum. However, because the Cherenkov angle is the complement of the angle of total-internal-reflection (TIR), emission from cascades which are downgoing, even only slightly, relative to the local surface tangent, suffer TIR and cannot be detected externally. This means that the Moon is an efficient filter that selects neutrinos over cosmic rays as the source particles for coherent radio emission. Similar effects of microwave TIR have, in fact, been observed in our proof-of-concept accelerator experiment.

The Goldstone experiment

To investigate the possibility of detecting pulsed, lunar-cascade radio emission, we have completed an initial series of dual antenna measurements using the Goldstone antennas DSS-14 and DSS-13. We expect a signal that is extremely broadband in nature, with a bandwidth that is roughly the inverse of the expected 1-ns pulse duration, on the order of 1 GHz. Since we are also sensitive to impulsive broadband terrestrial interference, we have developed a time-domain system that is akin to the fast pulse-counting methods employed at high-energy physics accelerators. In our case the “beam” is of cosmic origin, the “beam target” is the lunar regolith, and the detectors are DSN antennas.

To establish the lunar origin of any detected pulse we first require a <20-ns time coincidence between approximately equal band-limited pulses received in both left- and right-circular polarizations (LCPs and RCPs) at one antenna (DSS-14). The bandwidths employed here are 40–100 MHz—the full available bandwidth at S-Band (2.2 GHz). The requirement of equal LCP and RCP amplitude ensures that the incoming pulse is 100 percent linearly polarized (a necessary condition for the predicted coherent radiation from cascade emission).

Once such a highly polarized pulse is detected, the signal is recorded at a separate fre-

CONTINUED ON NEXT PAGE

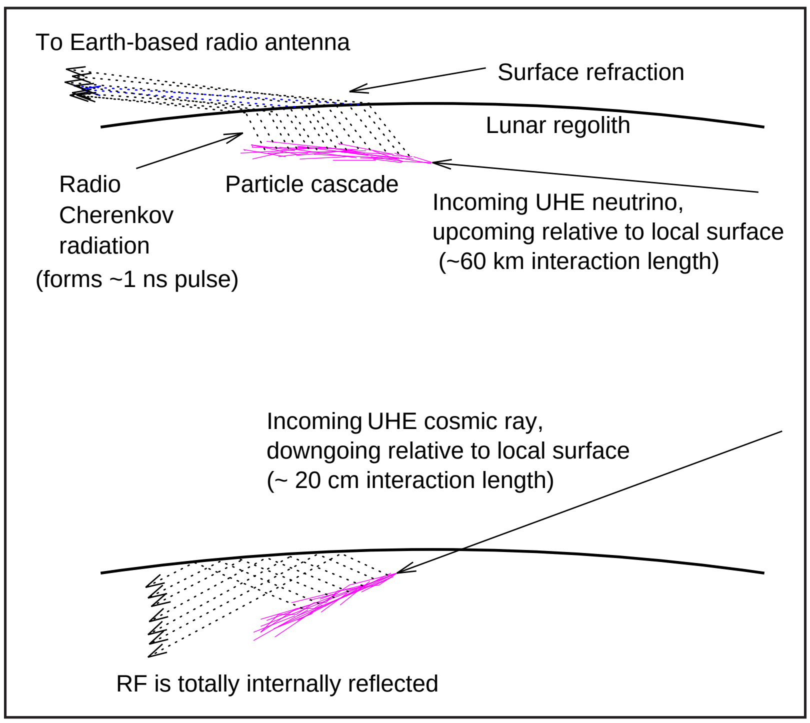

Figure 2. Geometry for neutrino and cosmic ray interactions in the lunar regolith

CONTINUED FROM PAGE 35

quency band (1.7–1.8 GHz; L-Band) whose feed is off-pointed in the sky so that it does not see the Moon. Since both the potential lunar pulses and any terrestrial interference are broadband, the slightly lower frequency receiver at L-Band will see any pulse that "flashes" across the sky (as interference does) or couples to the antenna in some other way. However, it does not receive any signal that originates from the Moon since it is off-pointed. This signal thus provides us with the ability to actively sense and reject radiofrequency interference signals.

Finally, a fourth signal from the 34-m antenna DSS-13, 21 km to the southeast, is also recorded. This signal comes from a bandpass that coincides with one of the polarizations used at DSS-14, and the antenna is directed to point at the Moon and track in a manner identical to DSS-14. Thus, if a coincidence trigger is formed at DSS-14, the data from DSS-13 is used to look for an additional pulse at an offset time that is appropriate to the range window of the lunar surface at the trigger time. This requirement provides a final stringent rejection of any terrestrial interference, and strong evidence that the pulse is of lunar origin.

In our initial test of these methods, we have observed the Moon for a total of 12 hours, spread over a number of different observing runs. Figure 3 shows an event exhibiting the characteristics of a pulse of lunar origin: the primary S-Band dual-circular polarization channels both show a large pulse, indicating a high degree of linear polarization. There is no corresponding pulse in the L-Band interference channel, indicating that local RF interference was not present. At the appropriate range gate (about 180 μs later in this case, determined by the lunar ephemeris and cross-checked with interferometric timing) in the data from the second antenna (bottom pane, with time axis shifted by geometric offset) there is also a pulse that appears to be well above the noise and is

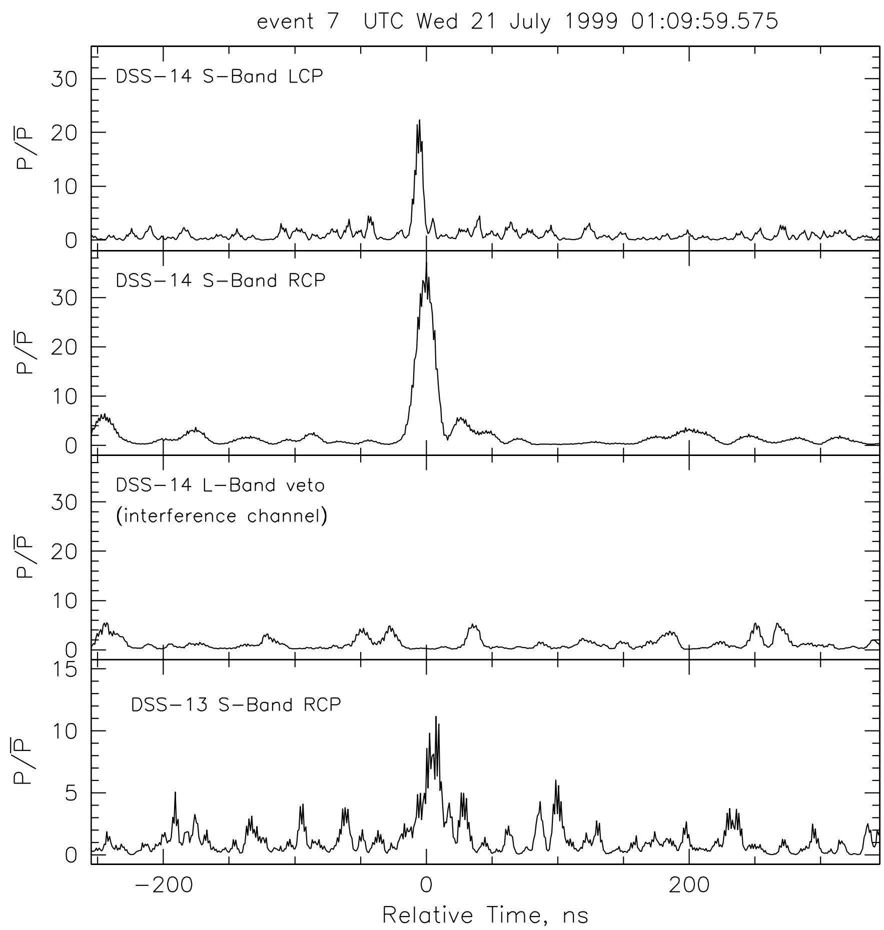

Figure 3. An example of a candidate pulse that may be of lunar origin. Each pane shows the plot of RF power divided by local thermal noise power. The top two panes show the two circular polarizations at the 70-m antenna, DSS-14, indicating a highly linearly polarized pulse has been detected. The third pane shows that there is no RF interference present. The bottom pane shows a comparable pulse at DSS-13, at the proper time-offset for the geometric delay from the Moon (here removed for clarity)

properly coincident with the pulse at DSS-14, suggesting that both may be of lunar origin.

While a handful of these intriguing events are present in our initial 12-hour data set, we have not yet fully assessed their statistical significance. Extraordinary claims certainly require extraordinary caution, and we are pursuing further observations to confirm the lunar origin of these events.

Sensitivity of the Approach

To indicate the possible importance of this approach, we have also estimated the effective flux sensitivity that our 12-hour proof-of-concept data implies, assuming that our understanding of the radio emission processes is correct. The results of this test are plotted in Figure 4, along with the predicted fluxes of high-energy neutrinos from a number of models, including AGN production, gamma-ray bursts, UHE cosmic-ray interactions, and three different early-universe topological defect models. Also plotted are the only existing limits at these energies, derived from about 70 days of livetime (requiring 2 years to achieve) of the ground-based cosmic-ray detector array known as the Fly's Eye, located in Dugway, Utah.

Our initial 90 percent confidence-level limit is shown plotted with inverted triangles. Even with our limited total integration time to date, we have already set the only limit in this energy regime, and have nearly achieved a sensitivity sufficient to constrain some of the theoretical flux predictions for UHE neutrinos. For comparison we also plot projected the Goldstone sensitivity level for a factor of 40 increase in livetime (20 days), which is within reach of our experiment over a period of several years. At present the sensitivity of this experiment appears to significantly exceed any other system in this energy regime, and further observations will continuously improve the limits.

Detection of unambiguous neutrino-induced cascades is still uncertain, but our goal is that the system will provide sensitivity to predicted phenomena in a region where few other instruments can presently search. Our proof-of-concept results have already generated significant interest among the particle astrophysics community, and have the potential to provide the first-ever detection of neutrinos from a cosmic source in this high-energy regime. We are on our way toward establishing the DSN as one of the world's first neutrino telescopes!

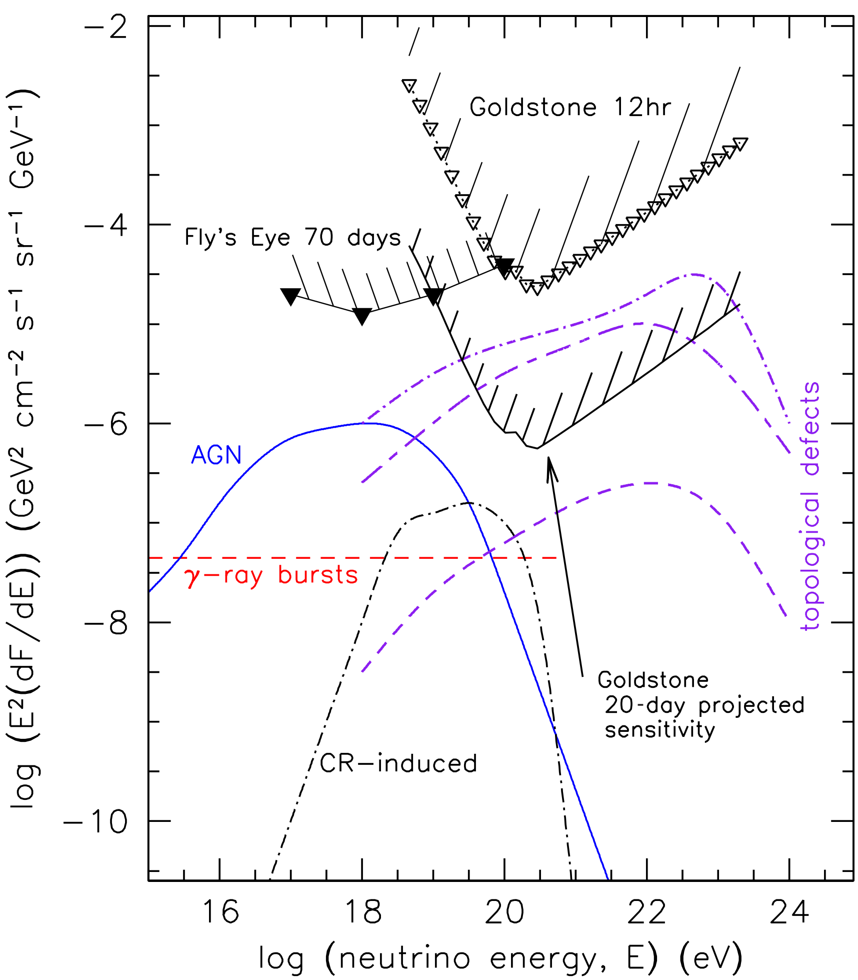

Figure 4. Predicted and experimental limits (both existing and projected) of fluxes of UHE neutrinos from various sources

MARS GLOBAL SURVEYOR AND THE DEEP SPACE NETWORK STUDY THE ATMOSPHERE OF THE RED PLANET

SAMI ASMAR AND TRISH PRIEST

INTRODUCTION

The Mars Global Surveyor (MGS) spacecraft was launched on November 7, 1996, and arrived at Mars on September 11, 1997. Over a six-month period after arrival, two sets of aerobraking maneuvers were performed to place the spacecraft into a two-hour circular orbit, with the goal of mapping the planet for one Martian year. Among the suite of MGS science instruments and investigations, the Radio Science investigation studies the atmosphere of the red planet and its gravitational field.

The atmosphere of Mars is studied via occultation measurements, where the spacecraft transmits a radio signal that travels through the atmosphere and is changed by it, prior to reception at the ground stations. The latter are equipped with specially designed receivers operating in an "open-loop" mode. A series of occultation measurements have been performed since the beginning of the mission, including the aerobraking phase. To date, nearly five thousand occultation measurements have been obtained (ingress and egress are counted separately). It is expected that by the completion of the prime mission phase, approximately eight thousand measurements will be made, with several thousand more possible during the likely extended mission. Compare this with two dozen occultations for the Galileo prime mission, for example.

Initial results from radio occultation measurements with MGS have yielded 88 vertical profiles of the neutral atmosphere. The Radio Science Team has reported [Hinson et al. 1999] that the measurements covered latitude 29 degrees North to 64 degrees South, mostly at nighttime during early summer in the Southern Hemisphere. Investigators have retrieved profiles of pressure and temperature versus radius and geopotential extending from the surface to the 10-Pa pressure level. The reported near-surface uncertainties in temperature and pressure are about 1 K and 2 Pa, respectively—far smaller than in previous radio occultations of Mars. The MGS Radio Science Team has also reported that the profiles resolve the radiative-convective boundary layer adjacent to the surface and reveal gravity waves. The data indicated the presence of a low-altitude westerly jet at southern summer solstice, as predicted by some general circulation models.

Figure 1 (from Dr. David Hinson, Stanford University) shows the temperature versus pressure and longitude at 66 degree North latitude on Mars, constructed from a set of 36 vertical profiles with roughly uniform spacing in longitude. The season was late spring and the local time of all measurements was about 0400. Contours of constant temperature are labeled in kelvins. The dashed line shows the pressure at the surface. This pattern of thermal structure remained stationary relative to

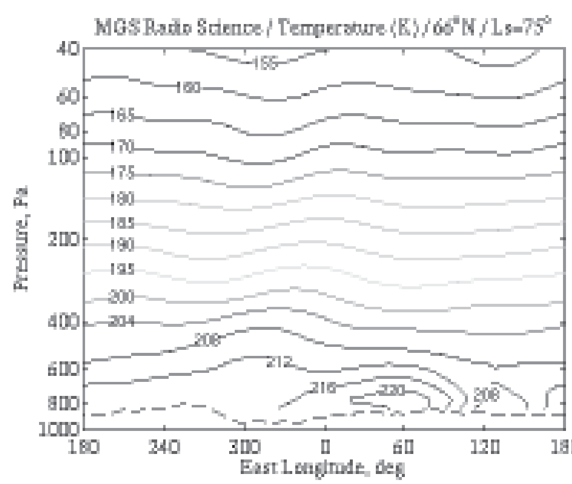

Figure 1. Temperature-pressure profiles from MGS Radio Science experiments

the surface during the one-week interval of the measurements. The largest temperature variations occur at low altitudes in the Eastern Hemisphere (0–180 degree E), where peak-to-peak variations on the 800-Pa pressure level exceed 12 K.

The Radio Science instrument is distributed between the spacecraft and the stations of the Deep Space Network (DSN). The elements of the instrument onboard the spacecraft are the X-Band telecommunications subsystem augmented by an ultra-stable oscillator (USO), with additional requirements placed on the attitude control subsystems. The ground components of the instrument include the Radio Science Subsystem as well as the Tracking Subsystem.

The central device of the DSN's Radio Science Subsystem is the Deep Space Communications Complex Spectrum Processor (DSP). As shown in Figure 2, the X-Band carrier signal from the spacecraft is downconverted via a predict-driven open-loop receiver that is controlled by the DSP. During a measurement, the DSP is directed to record approximately ten minutes of data without telemetry modulation, while the spacecraft's carrier signal is referenced to the onboard USO (one-way mode). This configuration allows for capturing the occultation egress and increases the power in the carrier.

In order to accomplish the challenging task of acquiring several thousands of radio occultation measurements, JPL's Radio Science Systems Group (RSSG) has designed and implemented a system of remote operations of the Radio Science instrumentation [Caetta et al. 1998] on behalf of the Deep Space Network. Engineers at JPL can directly operate the equipment. The system can

CONTINUED ON NEXT PAGE

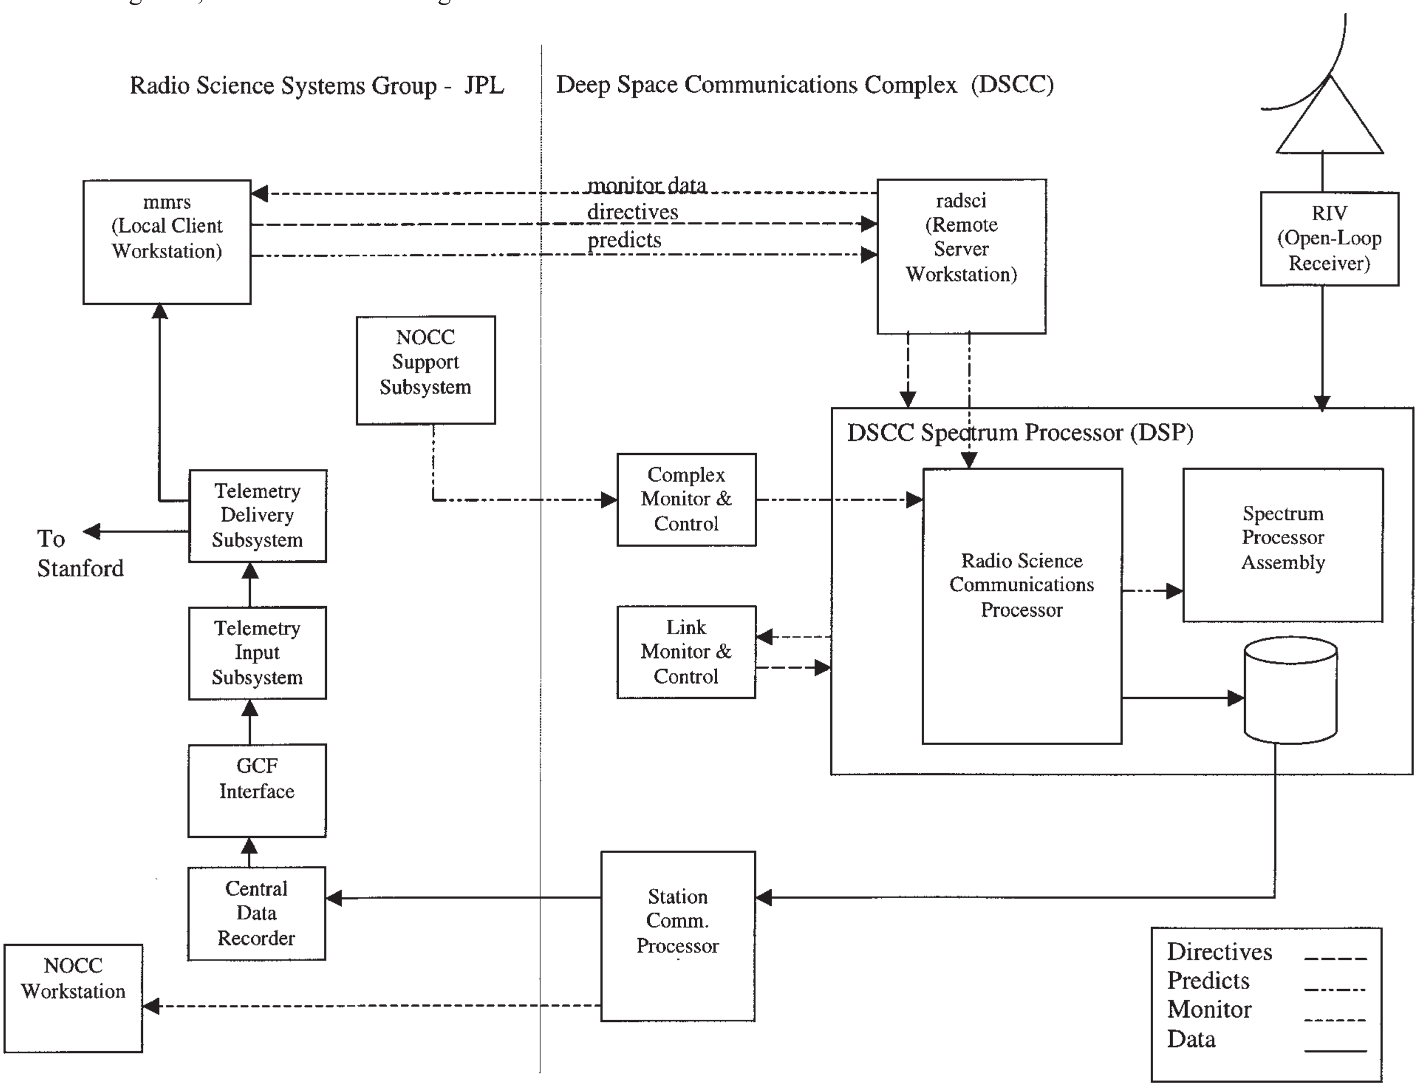

Figure 2. Block diagram of Radio Science remote operations

CONTINUED FROM PAGE 39

also be unattended for most occultation measurements by running scripted commands.

In the remote operations system, scripts are prepared to execute at specific times after the DSP has been placed in a link, or configured for a given spacecraft for a given pass. A workstation at JPL is used to set up scripts that run on similar workstations at each of the three DSN complexes. These workstations, in turn, send operator directives over a serial cable to the DSP, which is then automatically configured for each measurement by those scripts.

To date, the system has performed well, within 86 percent success rate for MGS occultation measurements. The remaining 13 percent were considered not useful for science analysis, with 1 percent considered degraded, mostly due to bad weather conditions. This excellent success rate can be compared with the only other large number of occultations, namely from the Pioneer Venus Orbiter, whose success rate over a 15-year mission was less than 50 percent.

The causes of MGS data loss are DSP failures and initialization. Another cause is filling the directory containing the frequency tuning predictions files, known as "predicts," leaving no space for additional files. In order to mitigate these problems, the first script of each DSN tracking pass is checked to ensure that the predict file, generated by the RSSG, is in place and that no DSP errors are detected.

In the process of improving the system, a new version of the remote operations software will detect and take appropriate action when an initialization failure occurs. The entire Radio Science Subsystem, however, is in the process of being replaced by a new digital instrument. The operational capabilities will be improved based on experience with the existing system, and more reliable hardware will be in place. These improvements will continue to make the DSN a world-class instrument for science research.

References

Caetta, J., S. Asmar, S. Abbate, M. Connally, and G. Golz, "Remote Operations of the Deep Space Network Radio Science Subsystem," TMO Progress Report 42-134, 1998.

Hinson, D. P., R. A. Simpson, J. D. Twicken, G. L. Tyler, and F. M. Flasar, "Initial results from radio occultation measurements with Mars Global Surveyor," *J. Geophys. Res.*, 104, E11, 26997–27012, 1999.

CONTINUED FROM PAGE 28

workforce requirements for future space operations. No matter how good your tools are, knowledge acquisition remains a fundamental bottleneck for the development of applications for BEAM and SHINE and for knowledge-base systems in general.

The ultimate goal of the DSN automation effort is "lights out" operation, which would achieve significant cost reduction. The lights, however, must be dimmed gradually. The DSN FDI automation work described in this paper has taken the first steps toward the dimming process. Another important aspect of our work is that of technology infusion into DSN operations. We have created a prototype that exercises innovative BEAM and SHINE technology in the conservative DSN operational environment. The ability to analyze all system data and perform system health assessment at all levels of operations is key to achieving autonomous reasoning and autonomous decision-making.

Acknowledgments

The work described was funded by the TMOD Technology Mission Services Program. Domain knowledge and data was provided by DSN expert Timothy Pham.

On the Web

The *TMOD Technology and Science Program News* and related TMOD features are located at:
http://tmot.jpl.nasa.gov/index.html
Click on "Program Overview Information," then click on "TMOD Program News."

Distribution

To have your name added to or deleted from the *TMOD Technology and Science Program News* distribution list, please call 4-9071.

The *TMOD Technology and Science Program News* is a publication of JPL's Telecommunications and Mission Operations Directorate (TMOD). The TMOD Technology Program is managed by James Lesh and the Science Program by Michael J. Klein.

Managing Editor .... Charles T. Stelzried
Associate Editor .... Michael Gregory

JPL
Jet Propulsion Laboratory
California Institute of Technology

JPL D-15493, Issue No. 12, 6/00